Review

A Review of Open-Source Frameworks for Real-Time Radar-Based Quantitative Precipitation Estimation: From Signal Correction to Operative Hydrological Validation

Florin Bîlbîe *

Posted Date: 24 April 2026

doi: 10.20944/preprints202604.1705.v1

Keywords: quantitative precipitation estimation; weather radar; open-source software; real-time processing; gauge-radar merging; hydrological validation; python frameworks

doi:10.20944/preprints202604.1705.v1

Review

A Review of Open-Source Frameworks for Real-Time Radar-Based Quantitative Precipitation Estimation: From Signal Correction to Operative Hydrological Validation

Florin Bîlbîe

Simion Mehedinti Doctoral School, University of Bucharest, Romania; florin.bilbie@s.unibuc.ro

Abstract

Real-time quantitative precipitation estimation (QPE) from weather radar is essential for hydrological forecasting, flash flood warning systems, and water resource management. Despite significant advances in radar technology and signal processing, operational QPE systems face persistent challenges including non-meteorological clutter contamination, signal attenuation, vertical profile biases, and systematic errors that require integration with ground-based rain gauge networks. This review synthesizes recent developments in open-source frameworks for radar QPE, spanning the complete processing chain from raw signal correction to operative hydrological validation. We examine state-of-the-art methods for clutter removal (polarimetric fuzzy logic, CLEAN-AP, neural network quality control), C-band attenuation correction (self-consistent and KDP-based approaches), and vertical profile of reflectivity (VPR) correction for warm-rain events. We compare gauge-radar merging techniques including mean field bias adjustment, spatially variable corrections, Kriging with External Drift (KED), and Conditional Merging, with emphasis on real-time applicability and look-back window strategies. The review identifies key open-source Python libraries (wradlib, Py-ART, pySTEPS, radproc, weatherDataHarmonizer) and documents operational latency constraints for flash flood warning systems. A critical research gap is identified: current open-source solutions lack documented workflows for integrating delayed 24-hour manual gauge readings into real-time QPE streams while maintaining low latency. This review provides researchers and practitioners with a comprehensive roadmap for developing robust, open-source, real-time radar QPE systems suitable for operational hydrological applications.

Keywords: quantitative precipitation estimation; weather radar; open-source software; real-time processing; gauge-radar merging; hydrological validation; python frameworks

---

1. Introduction

Quantitative precipitation estimation (QPE) from weather radar has become indispensable for operational hydrology, providing spatially continuous rainfall fields at high temporal resolution that are essential for flash flood forecasting, urban drainage management, and water resource planning. Unlike traditional rain gauge networks that offer point measurements with limited spatial coverage, weather radar systems can monitor precipitation over areas spanning hundreds of kilometers with spatial resolutions of 1-2 km and temporal updates every 5-15 minutes. However, the conversion of radar reflectivity measurements to accurate rainfall estimates remains challenging due to multiple sources of systematic and random errors.

The fundamental radar equation relates the received power to the radar reflectivity factor Z, which is then converted to rainfall rate R through empirical Z-R relationships. This seemingly straightforward process is complicated by numerous factors: non-meteorological echoes from ground clutter, buildings, and biological targets contaminate the signal; electromagnetic wave attenuation in

doi:10.20944/preprints202604.1705.v1

heavy precipitation leads to underestimation; the vertical profile of reflectivity introduces sampling biases when precipitation characteristics vary with height; and the inherent variability in raindrop size distributions makes Z-R relationships uncertain. These challenges are particularly acute for C-band radars (5.6 GHz), which are widely deployed in Europe and Asia but suffer from greater attenuation than longer-wavelength S-band systems.

To address these limitations, operational QPE systems integrate radar observations with ground-based rain gauge measurements through various merging techniques. Rain gauges provide accurate point measurements that serve as ground truth for bias correction and spatial adjustment of radar fields. However, gauge networks are often sparse, and many operational gauges report only daily accumulations with significant time delays, creating a mismatch with the real-time requirements of hydrological warning systems. The challenge of integrating delayed gauge data into real-time radar QPE streams while maintaining low latency remains a critical operational constraint.

The past decade has witnessed a proliferation of open-source software tools for radar data processing, driven by the scientific community's need for reproducible research and the operational community's demand for cost-effective, customizable solutions. Python has emerged as the dominant platform, with libraries such as wradlib [1], Py-ART [2], and pySTEPS [3] providing comprehensive functionality for radar data ingestion, quality control, correction algorithms, and nowcasting. These tools have democratized access to advanced radar processing techniques and enabled rapid prototyping of operational systems. However, the landscape of available tools is fragmented, and guidance on assembling complete end-to-end pipelines for real-time operational use remains limited.

This review addresses this gap by synthesizing recent advances in radar QPE methods with a focus on open-source implementations suitable for real-time operational deployment. We organize the review around three core pillars that correspond to the major stages of an operational QPE system: (1) radar signal quality corrections, including clutter removal, attenuation correction, and vertical profile adjustment; (2) gauge-radar integration methods, including bias correction and geostatistical merging; and (3) open-source software architecture and operational constraints, including computational performance and latency requirements for flash flood warning systems.

Our review is motivated by the practical needs of operational hydrologists and researchers developing real-time QPE systems. We emphasize methods that have been demonstrated in operational or near-operational contexts, document their performance characteristics, and identify the open-source tools that implement them. We pay particular attention to the temporal aspects of real-time processing, including look-back window strategies for bias correction and the integration of delayed gauge observations. Throughout, we highlight the trade-offs between accuracy, computational cost, and operational complexity that must be navigated in system design.

The structure of this paper follows the logical flow of an operational QPE system. Section 2 examines radar signal quality corrections, covering clutter removal, attenuation correction, and vertical profile adjustment. Section 3 reviews gauge-radar integration methods, comparing mean field bias correction, Kriging with External Drift, and Conditional Merging approaches. Section 4 surveys the open-source software ecosystem, documenting key libraries and their capabilities. Section 5 addresses operational constraints, including latency requirements and computational optimization strategies. Section 6 provides a comparative analysis and best practices framework for method selection. Section 7 identifies critical research gaps, with particular emphasis on the integration of delayed manual gauge data into real-time streams. Section 8 concludes with recommendations for future research and development.

2. Radar Signal Quality Corrections

The accuracy of radar-based QPE fundamentally depends on the quality of the input reflectivity measurements. Raw radar data are contaminated by multiple error sources that must be identified and corrected before rainfall estimation. This section reviews state-of-the-art methods for three critical quality correction steps: removal of non-meteorological clutter, correction of signal attenuation, and adjustment for vertical profile effects. We emphasize approaches that have been implemented in operational systems and are available in open-source frameworks.

2.1. Non-Meteorological Clutter Removal

Non-meteorological echoes, commonly referred to as clutter, arise from radar returns from stationary ground targets (buildings, terrain, wind turbines), moving biological targets (birds, insects), and atmospheric phenomena such as anomalous propagation. Clutter contamination is most severe at low elevation angles where the radar beam intersects the ground, yet these low elevations are preferred for QPE because they sample precipitation closest to the surface and minimize vertical profile biases. Effective clutter removal is therefore essential but must be balanced against the risk of removing legitimate precipitation echoes.

Modern clutter removal strategies leverage dual-polarization radar capabilities, which provide additional variables beyond reflectivity Z that characterize the physical properties of scatterers. The most widely used approach is **polarimetric fuzzy logic classification**, which combines multiple polarimetric variables—reflectivity (Z), differential reflectivity (ZDR), specific differential phase (KDP), and correlation coefficient (ρHV)—with texture fields and temperature information to classify each radar gate into hydrometeor types or clutter categories. Overeem et al. [1] document the implementation of polarimetric fuzzy logic in the Dutch operational real-time QPE system, where it serves as the first stage of quality control. The fuzzy logic approach assigns membership functions to each variable for different echo types and combines them through weighted aggregation to produce a classification confidence score. Gates classified as clutter with high confidence are removed or flagged for subsequent processing.

Complementing polarimetric classification, **CLEAN-AP** (Clutter Environment Analysis using Adaptive Processing) and related adaptive filtering techniques identify and suppress ground clutter through statistical analysis of temporal echo characteristics. CLEAN-AP exploits the fact that ground clutter echoes are statistically stationary over time scales of hours to days, whereas precipitation echoes are transient. By maintaining a clutter map derived from clear-air observations and applying adaptive thresholds, CLEAN-AP can remove persistent clutter while preserving precipitation signals. Husnoo et al. [2] demonstrate that CLEAN-AP, when integrated into a multi-stage quality control chain, significantly improves QPE accuracy on the UK weather radar network. Their study shows that combining CLEAN-AP with subsequent neural network-based quality control yields better performance than either method alone.

Neural network-assisted quality control represents a recent advance that addresses the limitations of rule-based approaches. Husnoo et al. [2] developed a neural network scheme that learns optimal elevation scan selection and quality control decisions from historical radar-gauge comparisons. The network is trained to minimize QPE errors by balancing the trade-off between low-elevation clutter contamination and high-elevation vertical profile biases. When integrated with CLEAN-AP preprocessing and vertical profile correction, the neural network QC scheme reduced mean absolute error by 15-20% compared to conventional rule-based approaches. The key advantage of the neural network approach is its ability to adapt to local conditions and learn complex, non-linear relationships between radar observables and QPE quality that are difficult to encode in explicit rules.

For operational implementation, **automated batch quality control workflows** are essential to process large volumes of radar data with minimal manual intervention. Kreklow [3] presents radproc, a GIS-compatible Python framework that facilitates automated preprocessing, quality control, and analysis of radar precipitation data. The radproc library provides tools for artefact detection, gauge data preparation, and integration into processing pipelines suitable for operational use. By automating routine QC tasks and providing standardized interfaces, such tools reduce the barrier to implementing sophisticated clutter removal algorithms in operational systems.

The practical recommendation emerging from the literature is to implement a **multi-stage clutter removal pipeline** that combines polarimetric fuzzy logic for initial classification, adaptive filtering (e.g., CLEAN-AP) for persistent clutter suppression, and optionally neural network-based scan selection to optimize the clutter-VPR trade-off. This layered approach provides robustness against diverse clutter types while maintaining computational efficiency suitable for real-time processing [1–3].

2.2. Attenuation Correction for C-Band Radar

Signal attenuation—the reduction in radar beam power as it propagates through precipitation—is a critical error source for C-band radars operating at 5.6 GHz. Attenuation causes systematic underestimation of reflectivity and rainfall rates, with errors increasing with path-integrated precipitation. For intense convective storms, uncorrected attenuation can lead to underestimation factors of 2-5 or more. Attenuation correction is therefore essential for accurate C-band QPE, particularly in regions prone to heavy precipitation.

The most widely adopted operational approach is the **self-consistent method**, which iteratively adjusts reflectivity and specific differential phase (KDP) to satisfy physical consistency constraints between attenuation and differential attenuation. The method exploits the relationship between specific attenuation A and KDP, both of which are proportional to the liquid water content along the propagation path. Gou et al. [4] present an improved self-consistent algorithm that incorporates additional constraints from polarimetric variables to stabilize the iterative solution and reduce sensitivity to measurement noise. Their method demonstrated reduced QPE errors in C-band observations compared to uncorrected data and simpler attenuation correction schemes. The self-consistent approach has the advantage of not requiring external calibration or assumptions about drop size distributions, making it robust across different precipitation regimes.

An alternative strategy is **KDP-based attenuation correction**, which directly estimates specific attenuation from KDP using empirical or theoretical relationships of the form $A = \alpha \cdot KDP^\beta$. Since KDP is immune to attenuation, calibration errors, and partial beam blockage, it provides a stable basis for attenuation estimation. Chen et al. [5] demonstrate the integration of KDP-based and specific attenuation (A-based) retrievals with traditional Z-based methods in a synthetic dual-band (S- and C-band) QPE system. Their results show that C-band A-based retrievals, after careful attenuation correction, can be less sensitive to raindrop size distribution variability than Z-based methods. The key challenge with KDP-based correction is accurate estimation of KDP itself, which requires careful filtering to remove noise while preserving the signal in regions of light precipitation where KDP is small.

Hybrid approaches that combine Z-based, KDP-based, and A-based rainfall estimators with seasonal or event-specific coefficients offer improved robustness. Chen et al. [5] show that integrating multiple polarimetric estimators and applying attenuation correction as a preprocessing step yields more consistent QPE across a range of precipitation intensities and types. The selection of estimator weights can be adapted based on precipitation regime, with KDP-based methods weighted more heavily in convective situations where attenuation is severe, and Z-based methods preferred in stratiform rain where KDP is noisy.

Despite the promise of machine learning for many radar processing tasks, the literature reviewed here does not provide demonstrated ML-based attenuation correction algorithms specifically validated for operational C-band QPE. This represents a potential area for future research, as ML approaches could potentially learn complex relationships between polarimetric observables and attenuation that are difficult to model explicitly.

For operational systems, the practical recommendation is to implement **improved self-consistent or KDP-anchored attenuation corrections** as a standard preprocessing step for C-band radar data. The choice between methods depends on data quality and computational constraints: self-consistent methods are more robust when KDP is noisy, while KDP-based methods are computationally simpler and perform well when high-quality KDP estimates are available [4,5].

doi:10.20944/preprints202604.1705.v1

2.3. Vertical Profile of Reflectivity Correction

The vertical profile of reflectivity (VPR) describes how radar reflectivity varies with height due to changes in precipitation microphysics, particularly the melting layer (bright band) and variations in raindrop size distributions. Since radar samples precipitation at heights ranging from a few hundred meters to several kilometers above ground (depending on range and elevation angle), and QPE requires surface rainfall estimates, VPR effects introduce systematic biases. These biases are particularly pronounced in warm-rain situations where reflectivity increases toward the surface due to coalescence and evaporation processes, and in stratiform precipitation with bright bands where reflectivity is enhanced in the melting layer.

Empirical VPR correction methods derive vertical profiles from observations and apply height-dependent adjustment factors to correct radar measurements to surface-equivalent values. Overeem et al. [6] developed a VPR correction algorithm specifically designed for warm-rain precipitation, which is common in maritime climates and was a dominant factor in the catastrophic July 2021 flooding in western Germany. Their algorithm derives VPR profiles from vertically pointing profilers or by analyzing the vertical structure in volumetric radar scans, then applies height- and range-dependent correction factors to the operational QPE product. In a case study of the July 2021 flood event, the VPR correction reduced normalized root mean square error by approximately 23% and normalized mean bias by approximately 20% compared to uncorrected radar estimates [6]. These substantial improvements demonstrate that VPR correction is not merely a refinement but an essential component of accurate QPE in warm-rain regimes.

The Dutch operational system, as documented by Overeem et al. [1], incorporates VPR correction as a standard processing step in their real-time gauge-adjusted radar precipitation product. The operational implementation uses a combination of climatological VPR profiles and real-time adjustments based on observed vertical structure. This hybrid approach balances the need for timely corrections (climatological profiles are always available) with the benefits of event-specific adjustments when observational data permit.

Neural network-based elevation selection offers an alternative approach to mitigating VPR effects. Rather than explicitly correcting for VPR, the neural network learns to select the optimal elevation angle or combination of elevations that minimizes QPE error for given conditions. Husnoo et al. [2] show that neural network QC, when combined with CLEAN-AP clutter removal and VPR correction, can reduce the trade-off between clutter contamination at low elevations and VPR-induced sampling errors at higher elevations. The network effectively learns a context-dependent strategy for balancing these competing error sources.

The effectiveness of VPR correction depends on the quality of the vertical profile information. In regions with dense radar networks or vertically pointing profilers, observed profiles can be used directly. In data-sparse regions, climatological profiles or profiles derived from numerical weather prediction models may be necessary. The key operational challenge is obtaining timely, representative VPR information with latency compatible with real-time QPE requirements.

For operational implementation, the recommendation is to **include VPR corrections derived from profilers, multi-elevation retrievals, or local VPR models**, particularly for warm-rain regimes and in regions with significant topographic relief. The substantial error reductions demonstrated in operational case studies justify the additional complexity [1,2,6].

3. Gauge-Radar Integration Methods

Despite advances in radar signal processing, systematic biases remain in radar QPE due to uncertainties in Z-R relationships, residual calibration errors, and unmodeled physical processes. Integration with ground-based rain gauge measurements is therefore essential to achieve the accuracy required for hydrological applications. This section reviews methods for gauge-radar integration, ranging from simple bias adjustment to sophisticated geostatistical merging techniques.

doi:10.20944/preprints202604.1705.v1

We emphasize approaches suitable for real-time operational systems and examine the temporal aspects of gauge data integration, including look-back window strategies.

3.1. Mean Field Bias and Spatially Variable Adjustments

Mean field bias (MFB) correction is the simplest and most widely used gauge-radar integration method. MFB computes a single multiplicative adjustment factor by comparing radar estimates with collocated gauge observations over a spatial domain and time window, then applies this factor uniformly to the entire radar field. The method is computationally efficient and requires minimal gauge density, making it attractive for operational systems. However, MFB assumes that radar bias is spatially uniform, which is often violated in practice due to spatially varying beam blockage, attenuation, and VPR effects.

Recognizing these limitations, operational systems have increasingly moved toward **spatially variable gauge adjustment** methods that allow the correction factor to vary across the radar domain. Overeem et al. [1] document the transition in the Dutch operational system from MFB to a spatially variable gauge adjustment implemented in 2023. The new method computes local adjustment factors using nearby gauges and applies spatial interpolation to produce a continuous correction field. This approach resulted in reduced average underestimation and improved capture of extreme precipitation events compared to the previous MFB-based system. The operational implementation uses a moving window approach where adjustment factors are updated at regular intervals (e.g., hourly) based on recent gauge-radar comparisons, allowing the system to adapt to evolving precipitation patterns and radar performance characteristics.

Moving median bias adjustment and related rolling-window techniques offer a middle ground between global MFB and fully spatial methods. Nielsen et al. [7] demonstrate the use of moving median adjustments to merge weather radar data with opportunistic rainfall sensors (personal weather stations, vehicle-mounted sensors). The moving median approach is robust to outliers and can accommodate heterogeneous sensor quality, making it suitable for integrating diverse gauge networks. The method computes adjustment factors over a sliding temporal window (e.g., the previous 24 hours) and applies them to current radar estimates, effectively implementing a look-back correction strategy.

Comparative evaluations consistently show that **bias-only adjustment methods are inferior to integration approaches** that combine radar and gauge information through statistical or geostatistical merging. Qingtai et al. [8] evaluated multiple merging methods in northern China and found that bias correction alone ranked worst among the tested approaches, while methods that explicitly integrate radar spatial structure with gauge point measurements performed best. This finding suggests that while bias adjustment is a necessary first step, it should be complemented by methods that leverage the complementary strengths of radar (spatial coverage) and gauges (point accuracy).

For operational systems with limited computational resources or sparse gauge networks, the practical recommendation is to implement **spatially variable gauge adjustment** rather than relying solely on global MFB. When gauge density permits, more sophisticated geostatistical merging methods (discussed in subsequent sections) should be considered [1,7,8].

3.2. Kriging with External Drift

Kriging with External Drift (KED), also known as regression kriging, is a geostatistical method that uses radar rainfall as a spatially continuous covariate (external drift) to guide the interpolation of gauge residuals. The method proceeds in two steps: first, a regression relationship between gauge observations and collocated radar estimates is established; second, the residuals (gauge minus predicted values) are interpolated using ordinary kriging and added back to the radar field. The result is a merged product that honors gauge observations exactly at gauge locations while using radar to inform the spatial structure between gauges.

KED has several attractive properties for operational QPE. It provides a statistically rigorous framework with quantifiable uncertainty (kriging variance), it naturally handles the different spatial supports of point gauges and gridded radar, and it can incorporate additional covariates such as elevation or distance to coast. Cassiraga et al. [9] demonstrate that **including temporal information and storm displacement** (Lagrangian coordinates) in the kriging framework improves interpolation skill. Their spatiotemporal KED approach uses the previous time interval and advection information to inform the covariance structure, effectively implementing a form of temporal persistence that is particularly valuable in real-time systems where recent observations provide strong constraints on current conditions.

Fast Bayesian regression kriging methods address the computational challenges of KED for large datasets. Yang and Ng [10] present algorithms that reduce the computational complexity of kriging through sparse approximations and efficient matrix factorizations, making KED feasible for operational systems with thousands of gauges and high-resolution radar grids. These computational advances are essential for real-time implementation, where processing must be completed within strict latency constraints.

The performance of KED depends critically on **gauge network density and variogram parameterization**. Shehu and Haberlandt [11] evaluated the relevance of merging radar and gauge data for rainfall nowcasting in urban hydrology and found that KED performs well with moderate to high gauge densities (>1 gauge per 100 km^2) and when variogram parameters are carefully calibrated. With sparse gauge networks, KED tends to revert toward the radar field in areas far from gauges, providing limited added value over bias-adjusted radar alone.

A key strength of KED is that it provides **kriging variance** as a measure of uncertainty, which can be propagated through hydrological models to quantify forecast uncertainty. This capability is particularly valuable for ensemble-based forecasting systems and risk-based decision-making frameworks.

For operational implementation, KED is recommended when gauge density is moderate to high, variogram parameters can be reliably estimated, and probabilistic uncertainty quantification is required. The method is particularly well-suited for applications requiring smooth, statistically grounded interpolation and explicit uncertainty estimates [9–11].

3.3. Conditional Merging and Hybrid Approaches

Conditional Merging (CM), also known as Double-Kernel Smoothing or Co-Kriging with Uncertain Means, takes a fundamentally different approach to gauge-radar integration. Rather than using radar as a drift variable for interpolating gauge residuals, CM treats both radar and gauge observations as uncertain estimates of the true precipitation field and produces a merged product that conditions radar spatial patterns on gauge totals. The method preserves the small-scale spatial variability of the radar field while adjusting intensities to match gauge observations.

The CM approach is particularly attractive for applications where **preserving radar spatial structure is critical**, such as nowcasting and hydrological modeling of small catchments where sub-grid precipitation variability drives runoff response. Melani [12] compared different radar-gauge merging methods for the Tuscany region and found that CM yields spatial patterns closer to raw radar while KED reflects gauge network structure. This difference has important implications for downstream applications: CM-merged fields better preserve the spatial organization of convective cells and frontal boundaries, which is essential for accurate nowcasting, while KED-merged fields provide smoother, more stable estimates suitable for climatological analysis.

Hybrid methods that combine elements of KED and CM have been proposed to balance their respective strengths. These approaches typically use KED to provide a smooth background field and CM to preserve high-frequency spatial variability, with the relative weighting determined by gauge density, radar quality, and application requirements. The flexibility of hybrid approaches allows them to adapt to varying data quality and operational constraints.

doi:10.20944/preprints202604.1705.v1

Comparative studies provide guidance on method selection. Shehu and Haberlandt [11] found that the choice between KED and CM depends on gauge density and intended use: KED is preferred for climatological applications and when gauge density is high, while CM is advantageous for nowcasting and when preserving radar spatial detail is critical. Qingtai et al. [8] confirmed that integration methods (both KED and CM) substantially outperform bias-only corrections, with the relative performance of KED versus CM depending on local conditions.

An important consideration for real-time systems is that CM requires **careful handling of radar and gauge uncertainties**. The method assumes that radar provides reliable spatial structure even if absolute intensities are biased, which may not hold in regions with severe beam blockage or attenuation. In such cases, hybrid approaches that downweight radar information in low-quality regions may be necessary.

For operational implementation, **Conditional Merging is recommended when preserving radar spatial detail is critical and gauge coverage is sparse**, particularly for nowcasting and small-catchment hydrology. KED is preferred when gauge density is higher and smooth, statistically grounded interpolation with uncertainty quantification is required [11,12].

3.4. Look-Back Window Strategies for Real-Time Systems

A critical challenge for operational QPE systems is the temporal mismatch between radar observations (available in near-real-time with latencies of 5-15 minutes) and gauge observations (often reported daily with delays of hours to days). This mismatch is particularly acute for manual rain gauges, which are typically read once per day at a fixed time (e.g., 08:00 local time) and may not be reported to operational systems until hours later. Integrating these delayed gauge observations into real-time radar QPE streams requires look-back window strategies that apply corrections retroactively or use historical gauge-radar relationships to adjust current estimates.

The **moving window bias adjustment** approach addresses this challenge by computing adjustment factors over a sliding temporal window (e.g., the previous 24-48 hours) and applying them to current radar estimates. The assumption is that systematic radar biases are temporally persistent over time scales of days, so that corrections derived from recent gauge-radar comparisons remain valid for current observations. Nielsen et al. [7] demonstrate this approach for merging radar with opportunistic sensors, showing that moving median adjustments computed over 24-hour windows effectively reduce bias while accommodating sensor reporting delays.

Overeem et al. [1] describe the operational implementation of look-back corrections in the Dutch real-time system. The system computes spatially variable adjustment factors using gauge observations from the previous 24 hours and applies them to current radar estimates. When new gauge observations become available (typically once per day for manual gauges), the adjustment factors are updated and the previous 24 hours of radar estimates are retroactively corrected. This approach provides both real-time estimates (using the most recent adjustment factors) and improved retrospective estimates (after gauge observations are incorporated).

Lagrangian persistence offers a more sophisticated approach to temporal integration. Cassiraga et al. [9] show that including storm displacement in the kriging framework improves interpolation skill by accounting for the advection of precipitation systems. In a look-back context, this means that gauge observations from upstream locations at earlier times can inform current estimates at downstream locations, effectively extending the temporal window over which gauge information is relevant.

The effectiveness of look-back strategies depends on the **temporal stability of radar biases**. Biases due to calibration drift, beam blockage, and VPR effects tend to be stable over time scales of days to weeks, making look-back corrections effective. However, biases due to attenuation and clutter contamination can vary rapidly with precipitation intensity and type, limiting the applicability of historical corrections to current conditions. Hybrid approaches that combine look-back bias adjustment with real-time quality control and attenuation correction offer the best performance.

A critical research gap identified in the literature is the lack of **documented workflows for integrating 24-hour manual gauge readings into real-time QPE streams**. While the literature describes methods for merging real-time automatic gauges and opportunistic sensors [7,11], and recommends inclusion of cross-border gauge data in operational products [1], no studies in the reviewed corpus provide detailed algorithms or implementation patterns specifically designed to ingest delayed manual gauge reports while maintaining low latency. This gap is particularly significant because manual gauges often provide the highest-quality precipitation measurements and are essential for long-term calibration and validation, yet their integration into real-time systems remains ad hoc.

For operational systems, the practical recommendation is to implement **rolling bias adjustments or spatially variable gauge adjustments combined with geostatistical merging**, using look-back windows of 24-48 hours to accommodate gauge reporting delays. When manual gauge observations become available, retroactive corrections should be applied to improve the quality of archived QPE products used for hydrological model calibration and climatological analysis [1,7,9].

4. Open-Source Software Ecosystem

The past decade has witnessed rapid growth in open-source software tools for radar data processing, driven by the scientific community's emphasis on reproducible research and the operational community's need for cost-effective, customizable solutions. Python has emerged as the dominant platform, offering a rich ecosystem of libraries for data ingestion, quality control, correction algorithms, visualization, and integration with hydrological models. This section surveys the key open-source tools available for building operational radar QPE systems, organized by functional category.

4.1. Core Radar Processing Libraries

wradlib (Weather Radar Library) is a comprehensive Python library for weather radar data processing, providing functionality for data I/O, georeferencing, quality control, attenuation correction, and rainfall estimation. wradlib supports multiple radar data formats (ODIM_H5, NEXRAD, Rainbow, etc.) and provides implementations of standard algorithms including beam blockage calculation, clutter filtering, attenuation correction (self-consistent and KDP-based methods), and gauge adjustment. The library emphasizes ease of use and educational value, with extensive documentation and tutorials that make it accessible to researchers and students. wradlib has been widely adopted in the research community and has been used in numerous operational and semi-operational systems.

Py-ART (Python ARM Radar Toolkit) was originally developed by the U.S. Department of Energy's Atmospheric Radiation Measurement (ARM) program for processing data from research radars, but has evolved into a general-purpose radar processing library. Py-ART provides robust data structures for radar data (the Radar object), comprehensive I/O capabilities, and implementations of advanced algorithms including dealiasing, attenuation correction, hydrometeor classification, and gridding. A key strength of Py-ART is its modular architecture, which allows users to easily extend functionality and integrate custom algorithms. Py-ART has strong support for polarimetric radar data and includes implementations of fuzzy logic classification and other quality control methods.

Both wradlib and Py-ART provide the foundational capabilities needed for radar data processing, but they differ in philosophy and design. wradlib emphasizes simplicity and educational value, while Py-ART emphasizes extensibility and integration with scientific workflows. In practice, many operational systems use both libraries, leveraging wradlib for standard processing tasks and Py-ART for specialized algorithms or when working with ARM radar data formats.

4.2. Nowcasting and Ensemble Generation

pySTEPS (Python Short-Term Ensemble Prediction System) is an open-source library for probabilistic precipitation nowcasting. Pulkkinen et al. [13] describe pySTEPS as a modular framework that implements multiple nowcasting methods, including deterministic extrapolation (Lagrangian persistence), stochastic nowcasting (STEPS), and ensemble generation. The library provides standardized interfaces for data I/O, optical flow estimation (for tracking precipitation motion), stochastic perturbation, and verification. pySTEPS has been designed with operational use in mind, offering computational efficiency and flexibility to accommodate different data sources and nowcasting algorithms.

A key feature of pySTEPS is its support for **ensemble nowcasting**, which quantifies forecast uncertainty by generating multiple equally likely realizations of future precipitation fields. Ensemble nowcasts are essential for probabilistic hydrological forecasting and risk-based decision-making. Imhoff et al. [14] demonstrate the operational application of pySTEPS for seamless rainfall and flood forecasting in the Netherlands, where pySTEPS ensemble nowcasts are blended with numerical weather prediction (NWP) model forecasts to provide continuous probabilistic precipitation forecasts from 0 to 48 hours. Their implementation shows that pySTEPS can meet operational latency requirements while providing skillful probabilistic forecasts.

The modular architecture of pySTEPS allows it to be integrated into broader operational workflows. The library can ingest quality-controlled radar data from wradlib or Py-ART, apply nowcasting algorithms, and export forecasts in formats suitable for hydrological models. This interoperability makes pySTEPS a valuable component of end-to-end operational systems [13,14].

4.3. Data Harmonization and Workflow Automation

radproc is a GIS-compatible Python framework for automated radar precipitation data processing, assessment, and analysis. Kreklow [3] designed radproc to facilitate large-scale, automated processing of radar data for climatological analysis and operational applications. The library provides tools for batch processing of radar composites, quality control, artefact detection, gauge data preparation, and export to GIS formats. A key strength of radproc is its emphasis on automation and reproducibility, with configuration-driven workflows that minimize manual intervention. radproc has been used for nation-scale radar data archives and operational QPE systems, demonstrating its scalability and robustness.

weatherDataHarmonizer addresses the challenge of integrating heterogeneous precipitation data sources (radar, nowcasts, NWP forecasts) into coherent time series suitable for hydrological models and warning systems. Grundmann and Akhmedova [15] describe weatherDataHarmonizer as a tool for harmonizing temporal resolution, spatial grids, and data formats across multiple precipitation products. The library handles the complexities of blending observations and forecasts with different latencies, update frequencies, and spatial resolutions, producing seamless precipitation time series that can be directly ingested by hydrological models. This capability is essential for operational flood forecasting systems that must integrate real-time radar observations with short-term nowcasts and longer-term NWP forecasts.

The combination of radproc for automated preprocessing and weatherDataHarmonizer for data integration provides a powerful foundation for operational QPE systems. These tools handle the "plumbing" of data ingestion, format conversion, and temporal alignment, allowing system developers to focus on algorithm implementation and hydrological modeling [3,15].

4.4. Parallel Processing and Computational Optimization

Real-time operational systems must process large volumes of radar data within strict latency constraints, requiring efficient computational implementations. Modern Python libraries increasingly leverage **just-in-time (JIT) compilation, multi-core parallelism, and GPU acceleration** to achieve the necessary performance.

doi:10.20944/preprints202604.1705.v1

Peer review discussions of the RoGeR (Runoff Generation Research) toolbox [16,17] highlight the benefits of JIT compilation (via Numba or JAX) and multi-backend support (CPU, GPU) for large-scale hydrological simulations. While RoGeR is a hydrological model rather than a radar processing tool, the computational strategies it employs are directly applicable to radar QPE pipelines. JIT compilation can accelerate computationally intensive algorithms such as attenuation correction, kriging, and ensemble generation by factors of 10-100 compared to pure Python implementations. GPU acceleration is particularly effective for embarrassingly parallel tasks such as processing multiple radar volumes, generating ensemble members, or computing spatial interpolations over large grids.

The pySTEPS library has incorporated computational optimizations including optional use of Numba for JIT compilation of critical loops and support for parallel processing of ensemble members. These optimizations enable pySTEPS to generate 10-member ensemble nowcasts with 10-minute time steps over domains of 1000×1000 km² in approximately 12 minutes on a 4-core workstation [14], meeting operational latency requirements for flash flood warning.

For operational systems, the practical recommendation is to **adopt JIT compilation and multi-core parallelism** for computationally intensive components of the processing chain. GPU acceleration should be considered for systems requiring very low latency or processing very large domains. The availability of high-level Python libraries (Numba, JAX, Dask) that provide these capabilities with minimal code changes makes optimization increasingly accessible to operational developers [14,16,17].

5. Operational Constraints and Real-Time Performance

Operational radar QPE systems for flash flood warning and hydrological forecasting must satisfy stringent latency requirements while maintaining high accuracy and reliability. This section examines the temporal constraints imposed by operational applications, identifies computational bottlenecks in typical processing pipelines, and reviews optimization strategies that enable real-time performance.

5.1. Latency Requirements for Flash Flood Warning

Flash floods are characterized by rapid onset, with time scales from rainfall initiation to peak discharge ranging from minutes to a few hours in small, steep catchments. Effective flash flood warning systems must therefore provide precipitation estimates and hydrological forecasts with latencies short enough to allow meaningful warning lead times. The total latency budget includes radar data acquisition and transmission, quality control and correction processing, gauge integration, hydrological model execution, and dissemination of warnings to end users.

Imhoff et al. [14] document operational latency measurements for a seamless rainfall and flood forecasting system in the Netherlands that integrates pySTEPS nowcasting with NWP forecasts and hydrological models. Their system generates 12-hour deterministic precipitation forecasts with 10-minute time steps in approximately 3.4 minutes on a 4-core workstation, and 10-member ensemble forecasts in approximately 12.3 minutes. Preprocessing (data ingestion, quality control, and regridding) requires approximately 10 minutes but is needed only when new NWP forecasts are issued (typically every 6 hours), so it does not contribute to the latency of every nowcast update. These measurements demonstrate that modern open-source tools can meet operational latency requirements for flash flood warning, provided that processing is properly optimized.

The latency budget can be decomposed into several components:

1. **Data acquisition and transmission** (1-5 minutes): Time for radar data to be collected, processed at the radar site, and transmitted to the central processing facility. This component is largely outside the control of the QPE system but can be minimized through efficient data formats and network infrastructure.

doi:10.20944/preprints202604.1705.v1

2. **Quality control and correction** (1-3 minutes): Clutter removal, attenuation correction, and VPR adjustment. These operations are computationally intensive but can be parallelized across radar volumes and elevation scans.
3. **Gauge integration** (0.5-2 minutes): Bias adjustment or geostatistical merging with gauge observations. Computational cost depends on gauge density and merging method, with simple bias adjustment requiring seconds and full KED requiring minutes for large domains.
4. **Nowcasting and ensemble generation** (3-12 minutes): Optical flow estimation, stochastic perturbation, and ensemble member generation. This is typically the most computationally expensive component for probabilistic systems.
5. **Hydrological model execution** (1-5 minutes): Distributed hydrological models for flash flood forecasting. Computational cost scales with catchment size, model complexity, and ensemble size.

For a typical operational system targeting 15-minute update cycles (matching radar volume scan intervals), the total processing latency should not exceed 10-12 minutes to allow time for data transmission and warning dissemination. This constraint is achievable with modern open-source tools and appropriate computational resources, as demonstrated by Imhoff et al. [14].

5.2. Computational Bottlenecks and Optimization Strategies

Profiling of operational radar QPE pipelines reveals several common computational bottlenecks:

Data I/O and regridding can consume a substantial fraction of processing time, particularly when working with large radar volumes in complex formats (e.g., ODIM_H5 with multiple elevation scans and polarimetric variables). Optimization strategies include: - Precomputing and caching regridding weights and coordinate transformations - Using efficient binary formats (HDF5, NetCDF4 with compression) for intermediate products - Implementing parallel I/O for multi-volume processing - Minimizing unnecessary data copies and format conversions

Quality control and attenuation correction involve computationally intensive operations such as fuzzy logic classification, iterative self-consistent attenuation correction, and KDP estimation. Optimization strategies include: - JIT compilation of critical loops using Numba or JAX - Vectorization of operations to exploit SIMD instructions - Parallel processing of elevation scans and azimuthal rays - Precomputation of lookup tables for fuzzy logic membership functions

Geostatistical merging (KED, CM) requires solving large linear systems or computing covariance matrices for thousands of gauge-radar pairs. Optimization strategies include: - Sparse matrix methods and efficient linear solvers - Local kriging approaches that limit the search radius - Fast approximations such as Fast Bayesian Regression Kriging [10] - Caching of variogram parameters and covariance structures

Ensemble generation in pySTEPS involves repeated application of stochastic perturbations and advection to generate multiple realizations. Optimization strategies include: - Parallel generation of ensemble members (embarrassingly parallel) - GPU acceleration of advection and perturbation operations - Efficient FFT implementations for spectral perturbation methods - Adaptive ensemble size based on forecast lead time

Kreklow [3] emphasizes the importance of **automated workflows and batch processing** for operational systems. By eliminating manual intervention and implementing robust error handling, automated pipelines can process large volumes of data reliably with minimal supervision. Configuration-driven workflows allow operational parameters to be adjusted without code changes, facilitating system tuning and adaptation to changing requirements.

The practical recommendation for operational systems is to **profile the complete processing chain to identify bottlenecks**, then apply targeted optimizations (JIT compilation, parallelization, algorithmic improvements) to the most expensive components. Modern Python tools make it feasible to achieve real-time performance on modest computational resources (4-8 core workstations) for domains up to 1000×1000 km² [3,14,16,17].

doi:10.20944/preprints202604.1705.v1

6. Comparative Analysis and Best Practices

The preceding sections have reviewed a wide range of methods for radar signal correction, gauge integration, and operational implementation. This section synthesizes these findings into a comparative framework and provides practical guidance for method selection based on application requirements, data availability, and operational constraints.

6.1. Method Selection Framework

Table 1 summarizes the key characteristics, strengths, limitations, and recommended applications for the major methods reviewed in this paper. The table is organized by processing stage (signal correction, gauge integration, software tools) to facilitate systematic method selection.

Table 1. Comparative Summary of Radar QPE Methods and Tools.

Category	Method/Tool	Key Strengths	Limitations	Recommended Applications	References
Clutter Removal	Polarimetric Fuzzy Logic	Robust classification using multiple variables; operational maturity	Requires dual-pol radar; parameter tuning needed	Standard first-stage QC for dual-pol systems	[1]
	CLEAN-AP	Effective for persistent ground clutter; adaptive	Requires clear-air reference; may miss transient clutter	Complement to polarimetric methods	[2]
	Neural Network QC	Learns optimal scan selection; adapts to local conditions	Requires training data; black-box nature	Systems with historical gauge-radar data	[2]
Attenuation Correction	Self-Consistent	No external calibration needed; robust across regimes	Computationally intensive; sensitive to noise	Standard for C-band operational systems	[4]
	KDP-Based	Computationally simple; immune to calibration errors	Requires high-quality KDP; noisy in light rain	When KDP quality is high; convective situations	[5]

VPR Correction	Empirical/Observed	Directly addresses local VPR effects; substantial error reduction	Requires profiler or multi-elevation data	Warm-rain regimes; regions with profilers	[1,6]
	NN Scan Selection	Optimizes clutter-VPR trade-off; adaptive	Requires training data	Systems with historical validation data	[2]
Gauge Integration	Mean Field Bias	Simple; minimal gauge requirements; fast	Assumes spatially uniform bias; limited accuracy	Sparse gauge networks; low-latency requirements	[1]
	Spatially Variable Adjustment	Accounts for spatial bias variation; operational maturity	Requires moderate gauge density	Operational systems with 1+ gauge per 200 km^2	[1,7]
	Kriging with External Drift	Statistically rigorous; provides uncertainty; smooth fields	Computationally expensive; requires variogram calibration	Moderate-high gauge density; uncertainty quantification needed	[9–11]
	Conditional Merging	Preserves radar spatial structure; good for nowcasting	Assumes radar structure is reliable; complex implementation	Nowcasting; small-catchment hydrology; sparse gauges	[11,12]
Software Tools	wradlib	Comprehensive; well-documented; educational	Less emphasis on operational optimization	Research; education; prototyping	-
	Py-ART	Modular; extensible; strong	Steeper learning curve	Advanced research; custom	-

	polarimetric support		algorithm development	
pySTEPS	Probabilistic nowcasting; ensemble generation; operational focus	Focused on nowcasting (not full QPE pipeline)	Nowcasting; probabilistic forecasting	[13,14]
radproc	Automated workflows; batch processing; GIS integration	Less emphasis on real-time algorithms	Large-scale archives; automated operational systems	[3]
weatherDataHarmonizer	Multi-source integration; temporal harmonization	Focused on data integration (not processing)	Blending radar, nowcast, NWP for hydrology	[15]

The selection of methods should be guided by a systematic assessment of:

1. **Data availability**: Dual-polarization radar data enable advanced clutter removal and attenuation correction; gauge network density determines the feasibility of geostatistical merging; availability of profiler or multi-elevation data enables VPR correction.
2. **Application requirements**: Nowcasting and small-catchment hydrology benefit from methods that preserve spatial structure (CM); climatological analysis and large-catchment hydrology benefit from smooth, statistically grounded methods (KED); flash flood warning requires low-latency methods.
3. **Computational resources**: Real-time systems with limited resources should prioritize computationally efficient methods (MFB, KDP-based attenuation correction); systems with more resources can employ sophisticated methods (self-consistent attenuation correction, KED, ensemble nowcasting).
4. **Operational maturity**: Operational systems should favor methods with demonstrated performance in similar contexts; research systems can explore newer methods (neural network QC, ML-based corrections).

6.2. Performance Metrics and Validation Approaches

Rigorous validation is essential to assess QPE system performance and guide method selection. Standard validation metrics include:

Point-to-pixel metrics compare radar estimates with collocated gauge observations: - Mean bias (MB) and normalized mean bias (NMB) quantify systematic over- or underestimation - Mean absolute error (MAE) and root mean square error (RMSE) quantify overall accuracy - Correlation coefficient (r) quantifies linear association - Categorical metrics (POD, FAR, CSI) assess detection of precipitation events above thresholds

Spatial verification metrics assess the spatial structure of precipitation fields: - Fractions Skill Score (FSS) evaluates spatial patterns at multiple scales - Structure-Amplitude-Location (SAL) metrics separately assess structural, amplitude, and location errors - Object-based verification identifies and tracks precipitation features

Hydrological validation assesses QPE performance for the intended application: - Streamflow simulation skill when QPE is used to force hydrological models - Flash flood warning skill metrics (lead time, false alarm rate, hit rate) - Reservoir inflow forecast accuracy

The literature reviewed here demonstrates the value of multi-metric validation. Overeem et al. [6] show that VPR correction reduces both RMSE (by ~23%) and normalized mean bias (by ~20%), indicating improvements in both random and systematic errors. Husnoo et al. [2] demonstrate that neural network QC reduces MAE by 15-20% compared to conventional methods. These quantitative performance improvements provide strong evidence for the value of advanced correction methods.

For operational systems, the practical recommendation is to implement **continuous validation** using independent gauge observations (not used in bias correction or merging) and to track performance metrics over time to detect degradation due to radar calibration drift, changing gauge networks, or evolving precipitation climatology. Automated validation workflows that generate daily or weekly performance reports enable proactive system maintenance [3].

7. Research Gaps and Future Directions

Despite substantial progress in radar QPE methods and open-source tools, several critical gaps remain that limit the accuracy, reliability, and operational utility of current systems. This section identifies the most significant gaps and outlines promising directions for future research and development.

7.1. Integration of Delayed Manual Gauge Data

The most critical gap identified in this review is the **lack of documented workflows for integrating delayed 24-hour manual gauge readings into real-time QPE streams**. Manual rain gauges, which are read once per day (typically at 08:00 local time) and may not be reported to operational systems until hours later, often provide the highest-quality precipitation measurements due to careful siting, regular maintenance, and quality control procedures. These gauges are essential for long-term calibration, validation, and climatological analysis. However, their integration into real-time systems remains ad hoc and poorly documented.

The reviewed literature describes methods for merging real-time automatic gauges [1,11], opportunistic sensors [7], and cross-border gauge data [1], but does not provide detailed algorithms or implementation patterns specifically designed to ingest delayed manual gauge reports while maintaining low latency. Key challenges include:

1. **Temporal alignment**: Manual gauges report 24-hour accumulations ending at a fixed time (e.g., 08:00), which must be disaggregated or aligned with radar time steps (typically 5-15 minutes).
2. **Retroactive correction**: When manual gauge observations become available (potentially 12-24 hours after the observation period), should the system retroactively correct archived radar estimates, or only update bias correction factors for future estimates?
3. **Quality control**: Manual gauge observations may contain errors due to transcription mistakes, gauge malfunctions, or observer errors. Automated QC procedures must detect and flag suspicious observations without discarding valuable data.
4. **Metadata management**: Manual gauge networks often have complex metadata (gauge relocations, changes in observation time, instrument changes) that must be tracked and incorporated into processing workflows.

Future research should develop and document **end-to-end workflows** for manual gauge integration that address these challenges. Promising approaches include: - Hybrid systems that use real-time automatic gauges for immediate bias correction and manual gauges for retrospective refinement - Temporal disaggregation methods that distribute 24-hour manual gauge totals across sub-daily time steps using radar temporal patterns - Automated QC procedures that flag suspicious manual gauge observations based on spatial consistency with nearby gauges and radar - Version-controlled QPE products that provide both real-time estimates (using available data) and retrospectively corrected estimates (after manual gauge integration)

7.2. Machine Learning for Attenuation Correction

While neural networks have been successfully applied to quality control and scan selection [2], the reviewed literature does not provide demonstrated **machine learning-based attenuation correction algorithms** specifically validated for operational C-band QPE. This represents a significant opportunity, as ML approaches could potentially learn complex relationships between polarimetric observables and attenuation that are difficult to model explicitly.

Promising research directions include: - **Deep learning models** that predict path-integrated attenuation from sequences of polarimetric observations along radar rays, trained on dual-frequency radar data or gauge-corrected rainfall estimates - **Physics-informed neural networks** that incorporate physical constraints (e.g., monotonic increase of attenuation with range, consistency between attenuation and differential attenuation) to improve generalization - **Transfer learning** approaches that adapt models trained on S-band or dual-frequency radar data to C-band systems with limited training data - **Uncertainty quantification** methods that provide confidence estimates for ML-based attenuation corrections, enabling adaptive blending with physics-based methods

The key challenge is obtaining sufficient high-quality training data, particularly for extreme precipitation events where attenuation is most severe. Dual-frequency radar observations, disdrometer measurements, and dense gauge networks can provide ground truth for model training and validation.

7.3. Heterogeneous Sensor Networks

The proliferation of **opportunistic precipitation sensors**—including personal weather stations, vehicle-mounted sensors, smartphone-based observations, and commercial microwave link networks—offers the potential to dramatically increase the spatial and temporal density of ground-truth observations. Nielsen et al. [7] demonstrate the value of integrating opportunistic sensors with radar and traditional gauges, but significant challenges remain:

1. **Quality control**: Opportunistic sensors have highly variable quality, with unknown calibration, siting issues, and potential for malicious or erroneous data. Robust QC procedures that can operate in real-time with minimal manual intervention are needed.
2. **Data assimilation**: Traditional gauge-radar merging methods assume homogeneous gauge quality and known measurement uncertainties. Methods that can optimally weight heterogeneous observations with varying quality are needed.
3. **Latency and availability**: Opportunistic sensors may have irregular reporting intervals and variable latency. Methods that can accommodate asynchronous, irregularly spaced observations are needed.
4. **Privacy and data access**: Many opportunistic sensor networks are operated by commercial entities or private individuals, raising questions about data access, privacy, and long-term availability.

Future research should develop **adaptive data assimilation frameworks** that can ingest heterogeneous sensor networks, automatically assess data quality, and optimally weight observations based on estimated uncertainty. Machine learning approaches that learn sensor-specific bias and error characteristics from historical comparisons with reference observations show promise for this application.

7.4. Seamless Integration with Hydrological Models

While this review has focused on radar QPE methods, the ultimate value of QPE is determined by its performance in hydrological applications. Current operational systems often have a disconnect

doi:10.20944/preprints202604.1705.v1

between QPE and hydrological modeling, with QPE products generated independently and then ingested by hydrological models. This separation can lead to suboptimal performance because QPE methods are not optimized for the specific requirements of hydrological applications.

Future research should explore **end-to-end optimization** of coupled QPE-hydrological systems, where QPE methods are tuned to maximize hydrological forecast skill rather than point-to-pixel agreement with gauges. Promising approaches include: - **Hydrologically consistent QPE** that preserves water balance at catchment scales - **Ensemble QPE** that provides spatially and temporally consistent uncertainty estimates for ensemble hydrological forecasting - **Adaptive QPE** that adjusts processing parameters based on hydrological model performance feedback - **Joint state estimation** that simultaneously estimates precipitation fields and hydrological states through data assimilation

The increasing availability of streamflow observations from operational networks and citizen science initiatives provides opportunities for hydrological validation and feedback that can drive QPE system improvements.

8. Conclusions

This review has synthesized recent advances in open-source frameworks for real-time radar-based quantitative precipitation estimation, spanning the complete processing chain from signal correction to operative hydrological validation. The key findings and recommendations are:

Signal Quality Corrections: Modern operational systems should implement multi-stage clutter removal combining polarimetric fuzzy logic, adaptive filtering (CLEAN-AP), and optionally neural network-based scan selection. For C-band radars, improved self-consistent or KDP-based attenuation corrections are essential and can reduce errors by factors of 2-5 in heavy precipitation. Vertical profile corrections, particularly for warm-rain events, provide substantial improvements (20-23% reduction in bias and RMSE) and should be included when profiler or multi-elevation data are available [1], [2,4–6].

Gauge-Radar Integration: Operational systems should move beyond simple mean field bias correction to spatially variable adjustments or geostatistical merging methods. Kriging with External Drift is recommended when gauge density is moderate to high (>1 gauge per 100 km²) and uncertainty quantification is required. Conditional Merging is preferred when preserving radar spatial structure is critical, particularly for nowcasting and small-catchment hydrology. Look-back window strategies using rolling bias adjustments over 24-48 hours can accommodate gauge reporting delays while maintaining real-time performance [1,7,9–12].

Open-Source Software: The Python ecosystem provides a comprehensive suite of tools for operational radar QPE. Core processing libraries (wradlib, Py-ART) handle data ingestion, quality control, and correction algorithms. pySTEPS enables probabilistic nowcasting and ensemble generation with operational latency. radproc and weatherDataHarmonizer provide automated workflows and data harmonization. Modern computational optimizations (JIT compilation, parallelization, GPU acceleration) enable real-time performance on modest hardware [3,13–15].

Operational Performance: Operational radar QPE systems can meet the latency requirements for flash flood warning (10-15 minute update cycles) using open-source tools on 4-8 core workstations. Deterministic nowcasts can be generated in 3-4 minutes, and 10-member ensemble nowcasts in 10-12 minutes. Careful attention to computational bottlenecks (I/O, attenuation correction, geostatistical merging) and targeted optimization is essential [14].

Critical Research Gap: The most significant gap identified in this review is the lack of documented workflows for integrating delayed 24-hour manual gauge readings into real-time QPE streams. Manual gauges provide the highest-quality precipitation measurements and are essential for long-term calibration and validation, yet their integration into real-time systems remains ad hoc. Future research should develop end-to-end workflows that address temporal alignment, retroactive correction, quality control, and metadata management for manual gauge integration.

Future Directions: Promising areas for future research include machine learning-based attenuation correction, integration of heterogeneous opportunistic sensor networks, and end-to-end optimization of coupled QPE-hydrological systems. The increasing availability of dual-polarization radar data, dense gauge networks, and computational resources creates opportunities for substantial improvements in operational QPE accuracy and reliability.

The open-source frameworks reviewed here provide a solid foundation for developing robust, accurate, real-time radar QPE systems suitable for operational hydrological applications. By combining state-of-the-art signal correction methods, sophisticated gauge integration techniques, and efficient computational implementations, operational systems can achieve the accuracy and latency required for flash flood warning and water resource management. Continued research to address the identified gaps, particularly manual gauge integration and ML-based corrections, will further enhance the operational utility of radar QPE systems.

Acknowledgments: The authors would like to acknowledge the use of SciSpace for AI-assisted literature analysis and drafting support. The final manuscript has been thoroughly reviewed and approved by the human authors.

References

1. A. Overeem, H. Holleman, and A. Buishand, “The Dutch real-time gauge-adjusted radar precipitation product,” *Earth System Science Data*, vol. 17, pp. 4715-4737, 2025. DOI: 10.5194/essd-17-4715-2025
2. N. Husnoo, T. Mitra, and D. Dufton, “A Neural Network Quality-Control Scheme for Improved Quantitative Precipitation Estimation Accuracy on the U.K. Weather Radar Network,” *Journal of Atmospheric and Oceanic Technology*, vol. 38, pp. 1741-1757, 2021. DOI: 10.1175/JTECH-D-20-0120.1
3. J. Kreklow, “Facilitating radar precipitation data processing, assessment and analysis: a GIS-compatible python approach,” *Journal of Hydroinformatics*, vol. 21, no. 4, pp. 652-670, 2019. DOI: 10.2166/HYDRO.2019.048
4. Y. Gou, H. Chen, and J. Zheng, “An improved self-consistent approach to attenuation correction for C-band polarimetric radar measurements,” *Atmospheric Research*, vol. 226, pp. 32-48, 2019. DOI: 10.1016/J.ATMOSRES.2019.03.006
5. J. Y. Chen, H.-C. Wang, and P.-L. Lin, “A Synthetic Quantitative Precipitation Estimation by Integrating S- and C-Band Dual-Polarization Radars,” *Remote Sensing*, vol. 13, no. 1, art. 154, 2021. DOI: 10.3390/RS13010154
6. A. Overeem, H. Holleman, and A. Buishand, “A Radar-Based Quantitative Precipitation Estimation Algorithm to Overcome the Impact of Vertical Gradients of Warm-Rain Precipitation: The Flood in Western Germany on 14 July 2021,” *Journal of Hydrometeorology*, vol. 24, pp. 1531-1552, 2023. DOI: 10.1175/jhm-d-22-0111.1
7. J. M. Nielsen, M. R. Rasmussen, and K. Thorndahl, “Merging weather radar data and opportunistic rainfall sensor data to enhance rainfall estimates,” *Atmospheric Research*, vol. 300, art. 107228, 2024. DOI: 10.1016/j.atmosres.2024.107228
8. Q. Qingtai, Z. Yong, and L. Changjun, “Evaluation of the radar QPE and rain gauge data merging methods in Northern China,” *Remote Sensing*, vol. 12, no. 3, art. 363, 2020. DOI: 10.3390/RS12030363
9. E. Cassiraga, N. Gómez-Hernández, and J. J. Gómez-Hernández, “Spatiotemporal Precipitation Estimation from Rain Gauges and Meteorological Radar Using Geostatistics,” *Mathematical Geosciences*, vol. 53, pp. 499-516, 2021. DOI: 10.1007/S11004-020-09882-1
10. P. Yang and T. L. Ng, “Fast Bayesian regression kriging method for real-time merging of radar, rain gauge, and crowdsourced rainfall data,” *Water Resources Research*, vol. 55, pp. 3194-3214, 2019. DOI: 10.1029/2018WR023857
11. B. Shehu and U. Haberlandt, “Relevance of merging radar and rainfall gauge data for rainfall nowcasting in urban hydrology,” *Journal of Hydrology*, vol. 594, art. 125931, 2021. DOI: 10.1016/J.JHYDROL.2020.125931
12. S. Melani, “Comparing different radar-raingauge precipitation merging methods for Tuscany region,” *EGU General Assembly 2023*, EGU23-12475, 2023. DOI: 10.5194/egusphere-egu23-12475

13. S. Pulkkinen, D. Nerini, A. Pérez Hortal, C. Velasco-Forero, U. Germann, A. Seed, and L. Foresti, "Pysteps: an open-source Python library for probabilistic precipitation nowcasting (v1.0)," *Geoscientific Model Development*, vol. 12, pp. 4185-4219, 2019. DOI: 10.5194/GMD-12-4185-2019
14. R. Imhoff, W. Klomp, R. Brauer, K. de Vries, C. Whan, and A. Overeem, "Towards seamless rainfall and flood forecasting in the Netherlands: improvements to and validation of blending in pysteps," *EGU General Assembly 2024*, EGU24-5849, 2024. DOI: 10.5194/egusphere-egu24-5849
15. J. Grundmann and N. Akhmedova, "Precipitation Data Harmonizer," *EGU General Assembly 2023*, EGU23-8978, 2023. DOI: 10.5194/egusphere-egu23-8978
16. Anonymous Reviewer, "Comment on 'RoGeR v3.0.3 – a process-based hydrological toolbox model in Python'," *Geoscientific Model Development Discussions*, 2023. DOI: 10.5194/gmd-2023-118-rc1
17. M. Steinbacher, P. Hövel, and T. Grünwald, "Reply to reviewer comments on 'RoGeR v3.0.3 – a process-based hydrological toolbox model in Python'," *Geoscientific Model Development Discussions*, 2024. DOI: 10.5194/gmd-2023-118-ac2