https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Process-aware mixture modelling of precipitation-event distributions using event-type clustering

Lena Ortega Menjivar[1], Nur Banu Özcelik[1], Johannes Laimighofer[1], and Gregor Laaha[1]

[1]BOKU University, Institute of Statistics

Correspondence: Ortega Menjivar, Lena (lena.ortega-menjivar@boku.ac.at)

Abstract. Accurate modelling of event-level precipitation distributions is critical for hydrological modelling, flood risk assessment, and climate studies, yet selecting suitable parametric families remains challenging. We propose a process-aware approach that first separates precipitation events into types and then constructs an overall mixture from the resulting type-specific distributions. Using 10-min precipitation records from two climatically contrasting Austrian weather stations (Graz Universität and Dornbirn, 2010-2022), we compare partitioning (k-means) with outcome-guided, model-based clustering implemented as finite mixtures of Gamma regressions (clusterwise distributional regression). Events are characterised by severity, magnitude, duration, intensity, time-to-peak, and lightning occurrence. While partitioning yields interpretable event groups, the outcome-guided approach provides substantially better distributional fit–reducing Kullback-Leibler divergence by 66-73% relative to a single-Gamma baseline–and improves visual diagnostics. Single-distribution baselines underrepresent process heterogeneity in precipitation; process-aware clusterwise mixtures that account for event-level structure address this gap. The framework is distribution-agnostic and portable, relying on information available at the event scale and lightning data, and can support applications such as hydrological modelling and stochastic weather generation.

1 Introduction

Fitting theoretical distributions to precipitation data plays a major role in many applications, such as hydrological modelling of rainfall-runoff (Beven, 2012, pp. 185–186), erosion studies (e.g. Petroselli et al., 2021), and climate modelling (e.g. Elshamy et al., 2006), often implicitly through rainfall generators that rely on precipitation distributions to predict non-zero, wet-day precipitation. Yet, despite their real-world importance, the question of *which* distribution to use to approximate wet-day precipitation is still very much open.

A well established choice is the two-parameter Gamma distribution, which has been used to model wet-day precipitation since at least Thom (1958), owing to its parsimony, positive support, and reasonably good fit to the bulk of precipitation data. It also found application in early stochastic rainfall generators such as WGEN (Richardson and Wright, 1984, p. 6). In their review of stochastic weather generators, Wilks and Wilby (1999) compared applications of the two-parameter Gamma, Exponential, and mixed Exponential distribution, and reported that the mixed Exponential distribution often outperforms the others, especially for heavy precipitation. They also mentioned non-parametric approaches where precipitation time series are resampled from existing data, but cautioned that these approaches may be less able to predict extreme events and may be

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

inflexible towards climate change. A range of other distributions have also been proposed, including the Weibull, Log-normal, mixed Log-normal, Kappa, generalised Beta, and Generalised Pareto distributions (Papalexiou and Koutsoyiannis, 2016).

In a more regionalised application, Li et al. (2013) evaluated the performance of Exponential, Gamma, Weibull, skewed Normal, mixed Exponential and hybrid Exponential/Pareto distributions in predicting precipitation time series for Canadian weather stations. In their comparative study, the mixed Exponential distribution performed best in simulating daily precipitation and hydrological modelling. In 2012, Papalexiou and Koutsoyiannis proposed the Generalised Gamma (GG) and Burr Type XII (a subtype of Generalised Beta) distributions as especially flexible three-parameter alternatives. In subsequent work (Papalexiou and Koutsoyiannis, 2013; Papalexiou et al., 2013; Papalexiou and Koutsoyiannis, 2016), they comparatively evaluated the Pareto, Log-normal, Weibull, Gamma, GEV, Gumbel, GG and BurrXII distributions across different seasons, spatio-temporal scales, and objectives. The overarching findings are that the Gamma and other light-tailed or exponential-type distributions systematically underestimate extreme rainfall. More flexible distributions (GG and BurrXII) are generally preferable, though the final choice is context-dependent.

The literature review makes clear that fitting distributions to wet-day (or sub-daily wet-interval) precipitation requires high flexibility regarding seasonal, temporal, regional, and meteorological patterns. The response of the hydrological community has accordingly been to introduce increasingly flexible and complex models.

We propose an alternative route to the necessary flexibility: explicitly accounting for process heterogeneity (Laaha et al., 2025). Precipitation totals arise from distinct meteorological processes (e.g., convective versus stratiform events), and a single distribution cannot, in general, accurately describe such heterogeneous event types. We argue that the lack of fit of commonly used distributions stems less from their inherent inflexibility than from ignoring event-level heterogeneity. Therefore, we emphasise focussing on *types of precipitation events* when fitting wet-day precipitation distributions. We achieve this by (1) defining precipitation events and their characteristics from a precipitation time series, (2) clustering similar events to identify groups, and (3) fitting theoretical distributions to each cluster, and then aggregating them into an overall distribution via weighted mixtures.

Other authors have used clustering of precipitation events successfully for their own specific research interests (e.g. Vásquez et al., 2024; Papacharalampous et al., 2025), which were mainly to gain descriptive insights into precipitation event types. However, to our knowledge, event clustering has not yet been systematically applied to improve the fit of theoretical wet-day precipitation distributions by deriving cluster-specific subdistributions.

Building on initial explorations in Laaha et al. (2025), the present work provides a complete workflow and frames the contribution in terms of research questions rather than distributional benchmarking. Specifically, we ask: (i) Does separating events into types and aggregating cluster-specific subdistributions into a finite mixture improve the fit of wet-interval precipitation distributions relative to a single-distribution baseline? (ii) Among alternative clustering strategies, how do partitioning clustering and outcome-guided, clusterwise distributional regression compare in terms of distributional fit and interpretability of event types? (iii) Where in the distribution (bulk versus tails) do improvements occur, and how do event characteristics (including lightning occurrence) relate to the observed gains?

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

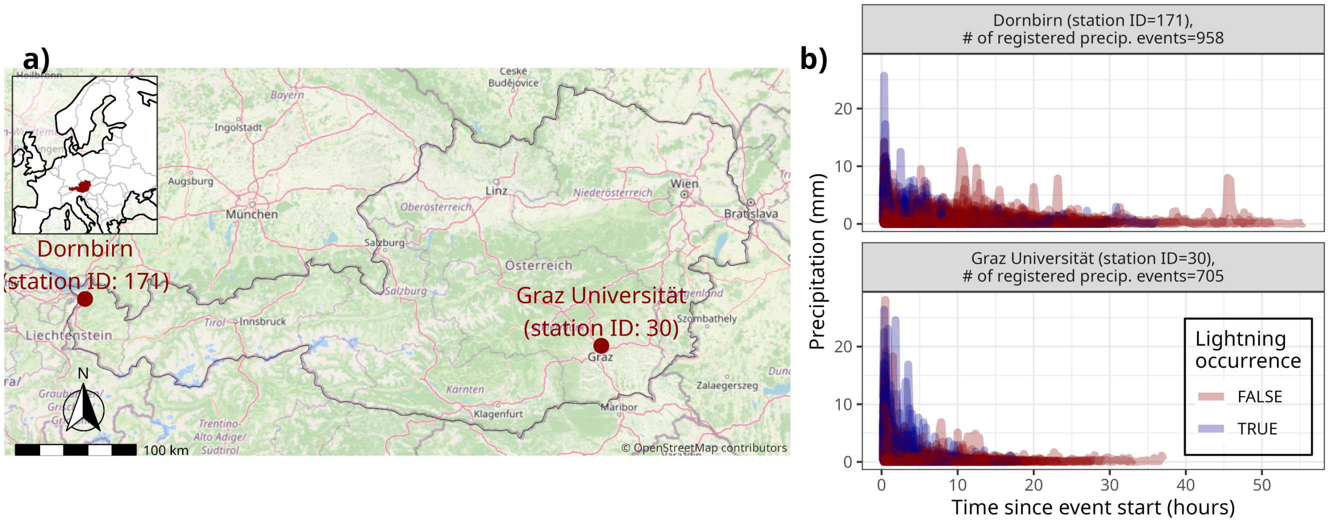

Figure 1. a) Location of the two stations (Graz Universität and Dornbirn) within Austria and Europe. b) Event-wise temporal distribution of precipitation for April–October, 2010–2022. Events are delineated from 10-minute data using a minimum intensity threshold of 0.2 mm (10 min)$^{-1}$ and a minimum inter-event time (MIT) of 4 h; colour indicates whether at least one lightning detection occurred within a 7 km radius of the station during the event window.

As a baseline for aggregation we choose the Gamma distribution as a simple, well-established example to isolate the value added by clustering; we do not benchmark against more flexible single-family models (e.g. BurrXII, GG), as that addresses a complementary question. Both aims are illustrated on 10-minute precipitation records from two Austrian weather stations with contrasting climatological regimes (Graz Universität and Dornbirn, 2010–2022).

2 Data and Methods

2.1 Data

We illustrate our proposed approach on two Austrian weather stations, Graz Universität and Dornbirn (Fig. 1a). Both are at similar altitudes (366 and 407.3 m a.s.l., respectively) but lie on opposite sides of the Austrian Alps – with a direct distance of approximately 435 km – and belong to precipitation regions with divergent diurnal precipitation patterns (Seibert et al., 2007; Yaqub et al., 2011). This setting allows us to test our approach under diverse Austrian conditions.

For these two stations, we used 10-minute precipitation data for the summer season (from April to October) of the years 2010 to 2022. The median recorded precipitation intensity is 0.2 mm (10 min)$^{-1}$ at both stations, and the maximum intensities are 25.7 and 28.1 mm (10 min)$^{-1}$, respectively.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

The precipitation data were obtained from GeoSphere Austria, https://data.hub.geosphere.at/dataset/klima-v1-10min (DOI: 10.60669/8fya-7x87).

To characterise generating processes, we included lightning as a proxy for convection events. Lightning data were provided by the Austrian Lightning Detection and Information System (ALDIS), which is part of the EUCLID (EUropean Cooperation for LIghtning Detection) network. The Lightning Detection System of ALDIS has an accuracy of 98% for downward (from cloud to the ground) flashes (Diendorfer et al., 2007). Data are available as the number of flashes per hour on a $1\,\mathrm{km} \times 1\,\mathrm{km}$ grid for Austria. Consistent and homogeneous data are available since 2010, whereas the median location accuracy is about 100 m for these years (Schulz, 2015). For each station, an event was flagged as convective if at least one flash occurred within a 7 km radius of the station during the event window. Within this radius, we counted 586 lightning occurrences in Dornbirn during our study period, and 963 in Graz.

2.2 Methods

Our proposed approach consists of four main steps:

1. Using merged precipitation and lightning time series, we define *precipitation events* by a transparent set of rules and derive *event-level characteristics* of interest.
2. After logarithmic transformation of skewed variables followed by min-max normalisation of all numeric characteristics, we cluster the precipitation events using two different approaches (partitioning and clusterwise regression) to enable methodological comparison. The resulting event-cluster assignments are mapped back to the time series by assigning each 10-minute observation to its event's cluster.
3. For each cluster in the precipitation time series, we fitted a Gamma distribution to describe the within-cluster rainfall distribution (using strictly positive 10-minute intensities).
4. We obtain an overall precipitation distribution as a weighted mixture of the cluster-wise distributions, with weights equal to the cluster share of 10-minute observations.

Finally, we compare our aggregated, cluster-wise mixture to a baseline Gamma distribution fitted to all positive observations irrespective of clusters. We evaluate the fits both visually and quantitatively via the Kullback-Leibler divergences D_{KL} (Kullback and Leibler, 1951) between the empirical and fitted theoretical distributions, using `R` (R Core Team, 2025) with packages `ggplot2` (Wickham, 2016) and `philentropy` (Drost, 2018). Detailed instructions for each step of the workflow follow:

2.2.1 Step 1: Defining precipitation events and their characteristics

To extract rainfall events from precipitation records, we set a minimum threshold of 0.2 mm $(10\ \mathrm{min})^{-1}$ precipitation and applied the classical minimum inter-event time (MIT) approach, merging sub-events separated by gaps shorter than the MIT (for early discussions see Todorovic and Zelenhasic, 1970). There is no universally accepted MIT value; rather, choices depend

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

on the research context. In their review, Dunkerley (2008) report MIT values ranging between 3 minutes and 24 h. We decided to set MIT to 4 h to distinguish convective events while maintaining the continuity of stratiform events. Additionally, we evaluated the robustness of this value by re-running our analyses with MIT values of 3, 6, 12, 18 and 24 h. For the robustness analysis scripts, and the resulting reports, see our supplemental Github repository (https://github.com/boku-stat/ClusteringRainfall/tree/v1.0.0, or alternatively https://doi.org/10.5281/zenodo.21358297).

For a MIT value of 4 h, we identified 958 precipitation events in the observed period for station Dornbirn, and 705 precipitation events in station Graz Universität (see Fig. 1b). Fig. 1b also provides a descriptive view of those events by showing the temporal distribution of precipitation for each event (with colour indicating whether lightning was registered during the event). In general, events in Dornbirn appear to be of lower intensity in each interval but longer duration than in Graz. This impression is confirmed when comparing the aggregated CDFs of event duration and 10-min level precipitation for the two stations via one-sided Kolmogorov-Smirnov tests ($\alpha = 0.05$).

To characterise our events, we follow Yevjevich's threshold-level approach (Yevjevich et al., 1967), which is widely used in hydrological drought studies (e.g. Tallaksen and Van Lanen, 2023; Laaha et al., 2017). Additionally, we calculate the time to the peak of the event to capture the temporal dynamics of precipitation events. The evaluated characteristics are:

- **Severity** S: total precipitation over the event ($S = \sum_{i=1}^{n} P_i$ [mm]),
- **Magnitude** M: maximum 10-minute precipitation within the event ($M = \max_{i=1,\dots,n} P_i$ [mm]),
- **Duration** D: event duration ($D = (t_{\text{end}} - t_{\text{start}}) + 10$ [min]),
- **Mean intensity** I: event-mean precipitation intensity ($I = (\sum_{i=1}^{n} P_i)/D$ [mm min^{-1}]),
- **Time to peak** T **(normalised)**: time from event start to the peak, normalised by duration ($T = ((t_{\max} - t_{\text{start}}) + 10)/D$ [-]),
- **Flash** F: presence of at least one lightning detection within a 7 km radius during the event window ($F \in \{0, 1\}$).

where

P_i is the precipitation in the i-th 10-minute interval [mm],

n is the number of 10-minute intervals in the event,

$t_{\text{start}}, t_{\text{end}}$ are the event start and end times (in minutes),

$t_{\max}$ is the time at which the highest 10-minute precipitation was recorded.

Summary statistics for these characteristics per station can be found in Appendix A.

2.2.2 Step 2: Clustering precipitation events

We comparatively evaluated two different clustering approaches regarding the interpretability of the derived event clusters; and the improvement of aggregated theoretical precipitation distribution fit:

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Clustering Approach I: Partitioning clustering

Partitioning methods are well established in hydrology. For example, Seibert et al. (2007) clustered Austrian weather stations regionally, and Papacharalampous et al. (2025) derived Alpine storm event types. Both studies used the Partitioning Around Medoids (PAM) algorithm. PAM was introduced by Kaufman and Rousseeuw (1990) for mixed-type data. While this is intuitive for our application (four numeric variables and one binary variable), we chose standard k-means for two reasons:

1. As Kaufman and Rousseeuw (1990, pp. 22–28) emphasise, one should distinguish between *symmetric* and *asymmetric* binary variables when selecting distance measures. Asymmetric binaries give different weight to 1 and 0; symmetric binaries weight both equally. In our case, both presence and absence are equally important for grouping precipitation events, so we used a distance suited to symmetric binaries—Simple Matching Distance and Euclidean distance are equivalent (up to a constant factor) for 0/1 variables (Kaufman and Rousseeuw, 1990, pp. 22–28).

2. Means provide the closed-form solution for clustering numeric variables under Euclidean distance and, for binary variables, can be interpreted as the cluster-wise marginal probability of observing a 1 (Leisch, 2006). Medoids are restricted to observed data points and, in moderately sized data sets, may yield suboptimal representatives of cluster centres.

We therefore performed k-means clustering based on Euclidean distances on the previously logarithmised and min-max scaled (apart from the binary lightning indicator) event characteristics using the *flexclust* functions (Leisch, 2006) with 10 reruns (default settings otherwise) and selected k by comparing cluster stability across 100 bootstrap samples via the Adjusted Rand Index (ARI; Hubert and Arabie 1985), using default settings.

Clustering Approach II: Clusterwise regression, outcome-guided learning approach

In contrast to Approach I, we used an "outcome-guided mixture-of-regressions framework" (*flexmix*; Grün and Leisch, 2007, 2025), which moves beyond purely unsupervised clustering. Building on our earlier work (Laaha et al., 2025), we modelled the within-event precipitation distribution as the response and let event characteristics serve as predictors.

Conceptually, clusters are defined by the regression model that best predicts each event's precipitation distribution, rather than by feature similarity. Events are assigned to the cluster whose common model most accurately links magnitude, duration, time-to-peak, and lightning to 10-minute precipitation levels.

The overall model is a finite mixture:

$$f(\mathbf{y}_i \mid \mathbf{x}_i) = \sum_{k=1}^{K} \pi_k \prod_{j=1}^{n_i} f_k(y_{ij} \mid \mathbf{x}_i, \boldsymbol{\theta}_k), \qquad \pi_k \geq 0, \; \sum_{k=1}^{K} \pi_k = 1 \tag{1}$$

where $\mathbf{y}_i = (y_{i1}, \ldots, y_{in_i})$ is the vector of 10 minute precipitation observations for event i, n_i is the number of observations in event i, and $\mathbf{x}_i$ contains the event characteristics. The π_k are the mixing proportions, $f_k(\cdot)$ is the cluster-specific density, and $\boldsymbol{\theta}_k$ are cluster-specific parameters. The product over j reflects the grouping constraint - all n_i observations of event i share the same cluster membership.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

For each cluster k, we fitted a Gamma regression model to the precipitation intensities (adding a small constant to ensure strictly positive values), using a log-link with main effects of magnitude, duration, time-to-peak, and lightning, plus interactions between lightning and the other characteristics to allow for distinct thunderstorm dynamics. We selected this subset of the derived event characteristics based on exploratory analyses, in which it consistently yielded non-degenerate solutions and stable convergence of the EM algorithm. Parameters are estimated using the Expectation-Maximization (EM) algorithm, which iteratively assigns events to clusters and updates regression coefficients until convergence (maximum 100 iterations). As in the partitioning approach, we selected the number of components k by assessing clustering stability across bootstrap resamples using the Adjusted Rand Index (ARI).

Chosen numbers of clusters *k*

We selected the number of clusters k for each model-station-combination via bootstrapping the data 100×2 times, applying the respective clustering algorithm, and calculating pairwise ARIs to assess stability (see Leisch, 2006 for methodological details).

For the partitioning approach, this procedure indicated $k = 3$ for both stations, based on the "elbow" in the median ARI values. In the outcome-guided mixture-of-regressions approach, we selected

- $k = 4$ for station Dornbirn (station ID = 171), as it corresponds to the "elbow" median ARI, and
- $k = 2$ for station Graz (station ID = 30), as it yields the highest median ARI and no clear "elbow" is present.

Visualisations of the ARI distributions are provided in Appendix B.

2.2.3 Step 3: Fitting Gamma distributions to precipitation data

We fitted Gamma distributions to strictly positive precipitation intensities, focusing on wet-interval (wet-day) distributions. Specifically, we excluded zeros and sub-threshold values (at or below the wet-interval threshold used in event definition) and estimated parameters by L-moments—once using all positive precipitation observations to obtain a baseline distribution, and once separately for each cluster. We used a shape-rate parametrisation (α,β), with density:

$$f(y \mid \alpha, \beta) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} y^{\alpha-1} e^{-\beta y}, \qquad y > 0,\ \alpha, \beta > 0. \tag{2}$$

Here, $\Gamma(\cdot)$ denotes the Gamma function. L-moment estimation is used because it is more robust to outliers than maximum likelihood for skewed precipitation data (Hosking, 1990, 2024; Zaghloul et al., 2020).

2.2.4 Step 4: Deriving a cluster-aggregated distribution

To obtain a single overall distribution that accounts for cluster structure, we combine the cluster-specific Gamma distributions as a finite mixture weighted by relative cluster sizes:

$$f_{\text{agg}}(y) = \sum_{k=1}^{K} w_k\, f_k(y \mid \alpha_k, \beta_k), \qquad w_k = \frac{n_k}{\sum_{k=1}^{K} n_k}. \tag{3}$$

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Here, w_k is the proportion of precipitation observations assigned to cluster k, n_k is the number of observations in cluster k, and f_k is the corresponding cluster-specific Gamma density with shape–rate parameterisation (α_k, β_k). We then compare this cluster-aggregated distribution with the baseline (single) Gamma fit to assess whether accounting for cluster structure improves the representation of the precipitation distribution.

3 Results

3.1 Exploratory analysis of precipitation event clusters

We begin with an exploratory view of event characteristics and assigned partitions via Factor Analysis for Mixed Data (FAMD), implemented in the `R` package `FactoMineR` (Lê et al., 2008).

Fig. 2 shows FAMD plots for the two study stations; the left panels display clusters from the partitioning approach, and the right panels display clusters from the clusterwise regression approach. Apart from cluster membership (colour), all other plot elements are identical. Dimensions (Dim.) 1 and 2 together account for 65.3% of total inertia at Graz, and 62.0% at station Dornbirn, the first four dimensions collectively explain 93.5% and 92.0% respectively. For both stations, the first two dimensions are primarily aligned with event `severity` and `magnitude` (29.0% and 34.2% contribution to Dim. 1 at Graz; and 34.6% and 31.4% at Dornbirn) and event `duration` (47.0% contribution to Dim. 2 at Graz; and 34.2% at Dornbirn), with duration playing a comparatively larger role in Graz. Dimension 3 focuses on `time-to-peak` (47.0% at Graz; 72.8% at Dornbirn), whose influence is clearly stronger in Dornbirn. Dimension 4 most strongly reflects lightning occurrence (`flash`: 54.6% at Graz, 59.5% at Dornbirn). On Dim. 4 we also observe similarities between `flash` and `time_to_peak` at Graz; `time_to_peak` contributes more strongly to this dimension at Graz (36.3%) than at Dornbirn (23.3%).

These dimensional patterns point to distinct event structures at the two stations, which we revisit in the Discussion (Sect. 4) in light of their differing meteorological settings.

Regarding the event clusterings, the partitioning clustering approach (left panels) yields clear separations along the axes of the event characteristics, with near-perfect separation by `flash`. For the clusterings obtained via clusterwise regression, the clusters are harder to interpret using only the depicted event characteristics. While the groups roughly follow the axes of the original variables (with an especially clear separation by event `magnitude`), there is much more overlap between clusters. Notably, the separation by `flash` is much less clear, as clusters contain events both with and without lightning. By contrast, the clusterwise regression approach incorporates the event-level precipitation distribution alongside the characteristics shown in the FAMD plots (see Subsection 2.2.2). This indicates that event-level precipitation distributions are shaped by factors beyond the characteristics analysed here.

This interpretation is supported by the comparison of cluster-wise Empirical Cumulative Distribution Functions (ECDFs) of precipitation in Fig. 3. For the partitioning approach, the ECDFs for the first two clusters closely match the overall ECDF at both stations; the only divergent group is the third cluster (the "lightning cluster"). Under the clusterwise regression approach, the ECDFs show much clearer differentiation. Incorporating the event-level precipitation distribution into the clustering yields

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

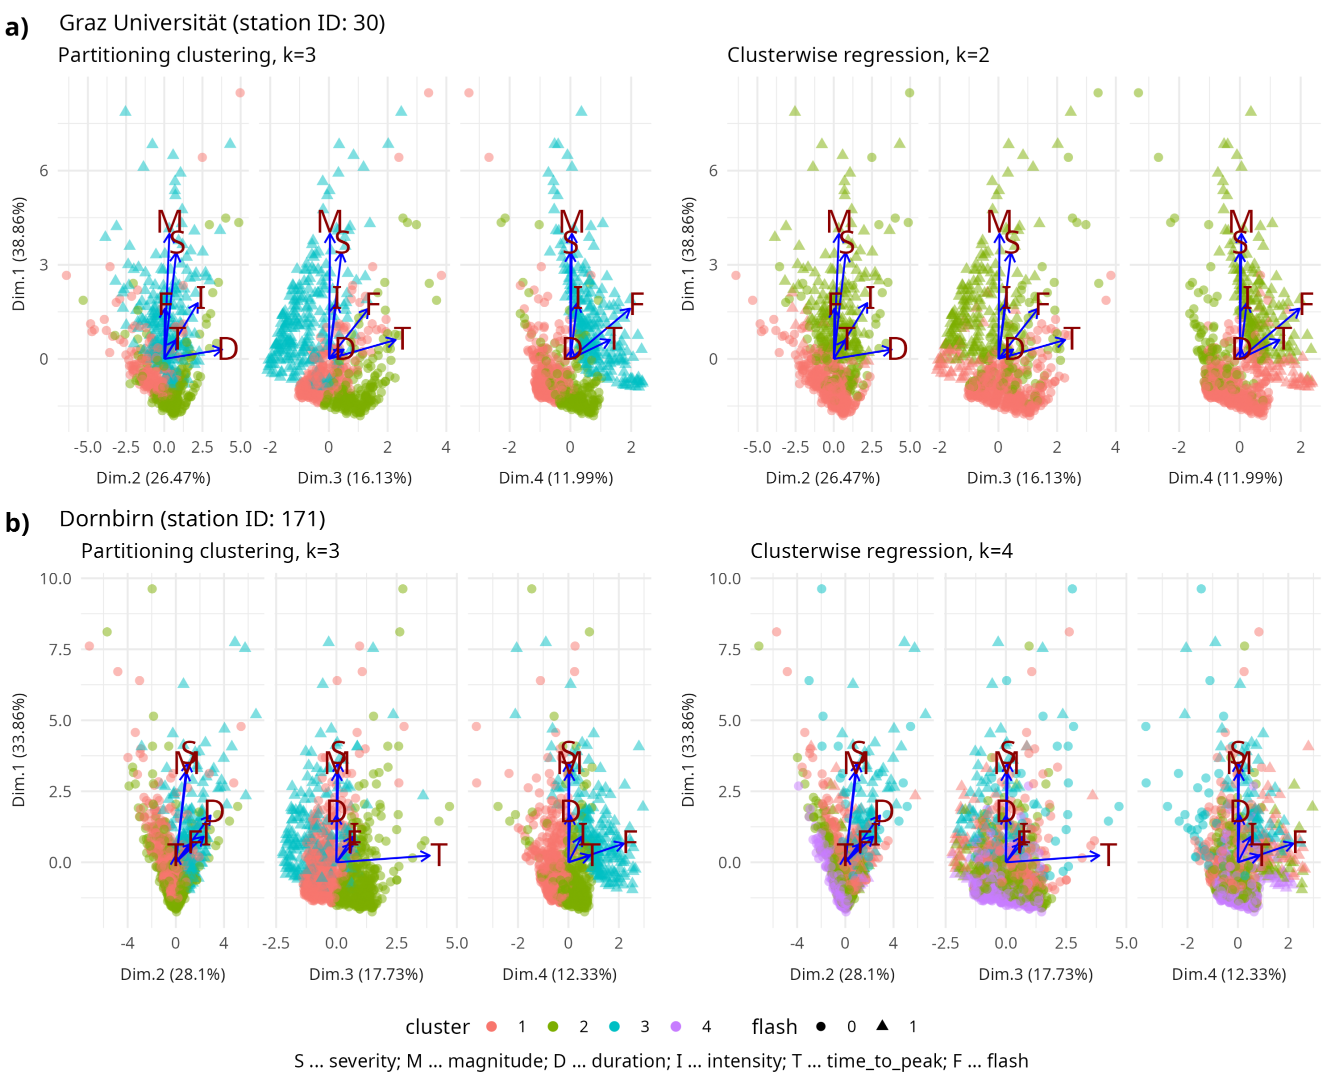

Figure 2. FAMD plots for event characteristics at two stations: a) Graz Universität and b) Dornbirn. Each panel contrasts Dimension 1 (largest component) with Dimensions 2-4. Percentages on the axis labels indicate the variance explained by each dimension. Arrows show variable contributions (loadings); points are events projected into FAMD space. Colour denotes cluster membership, shape indicates lightning occurrence.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

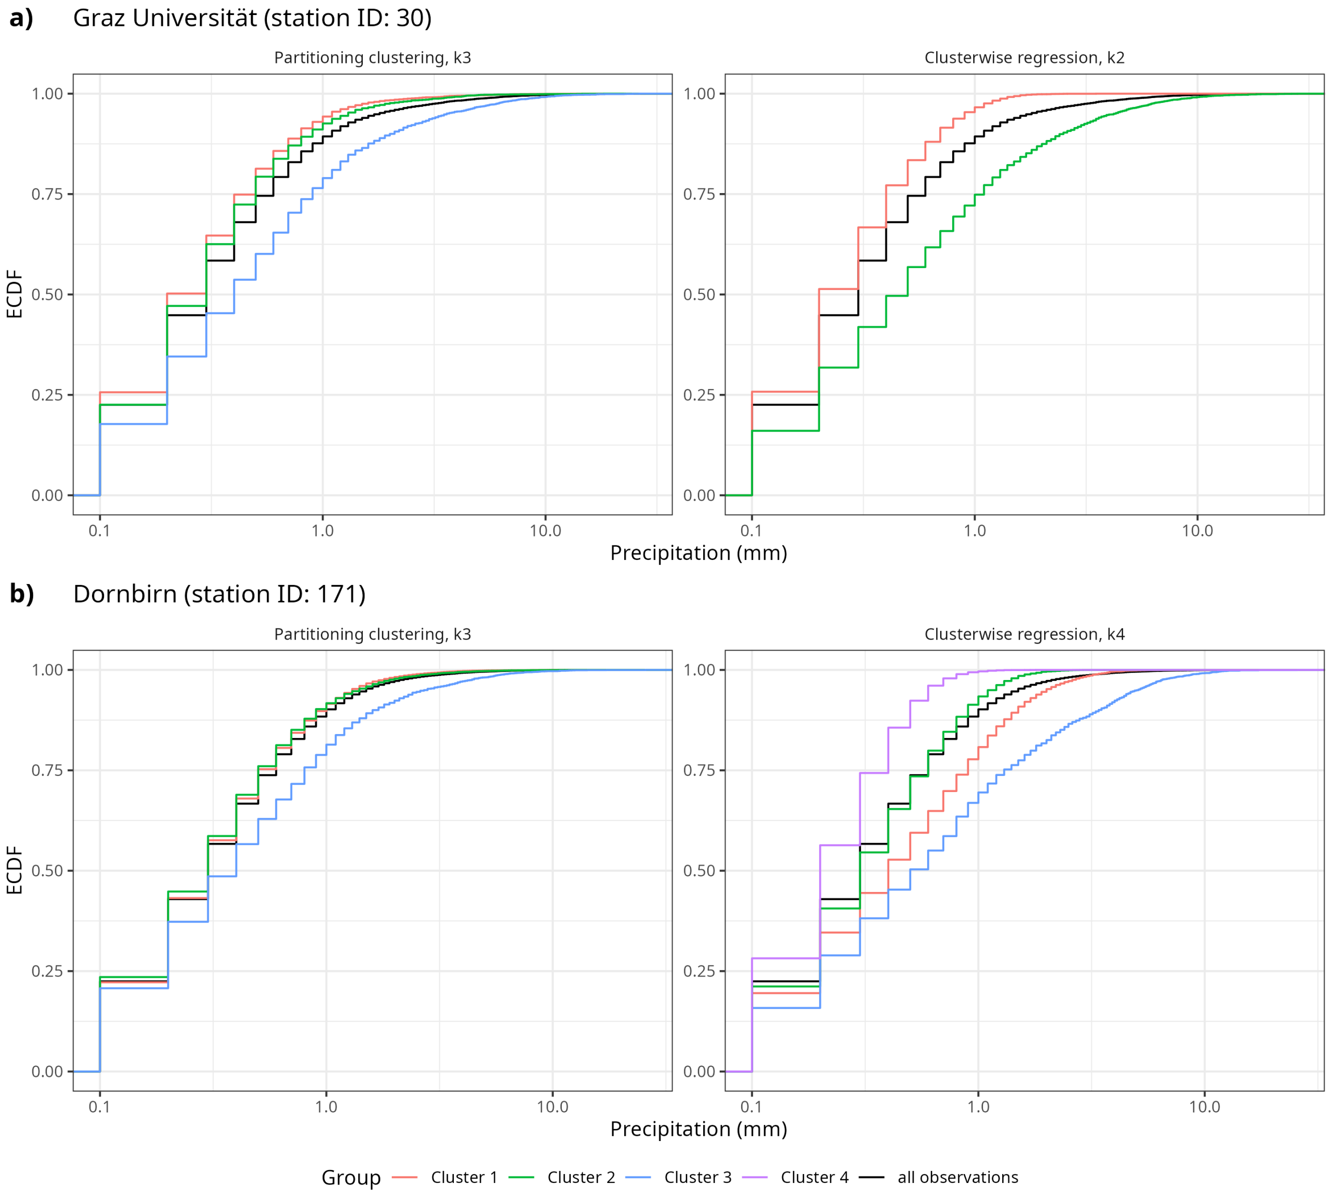

Figure 3. Empirical Cumulative Distribution Functions (ECDFs) of precipitation for both stations (x-axis on a logarithmic scale). Black line: overall ECDF; Coloured lines: clusterwise ECDFs. Columns compare the overall ECDF with ECDFs from (left) partitioning clustering and (right) clusterwise regression. Rows show (top) Graz Universität and (bottom) Dornbirn.

greater separation in the overall precipitation distribution—one of the key results of this study. Whether this separation translates into improved modelling of the rainfall distribution is assessed in the subsequent analysis.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

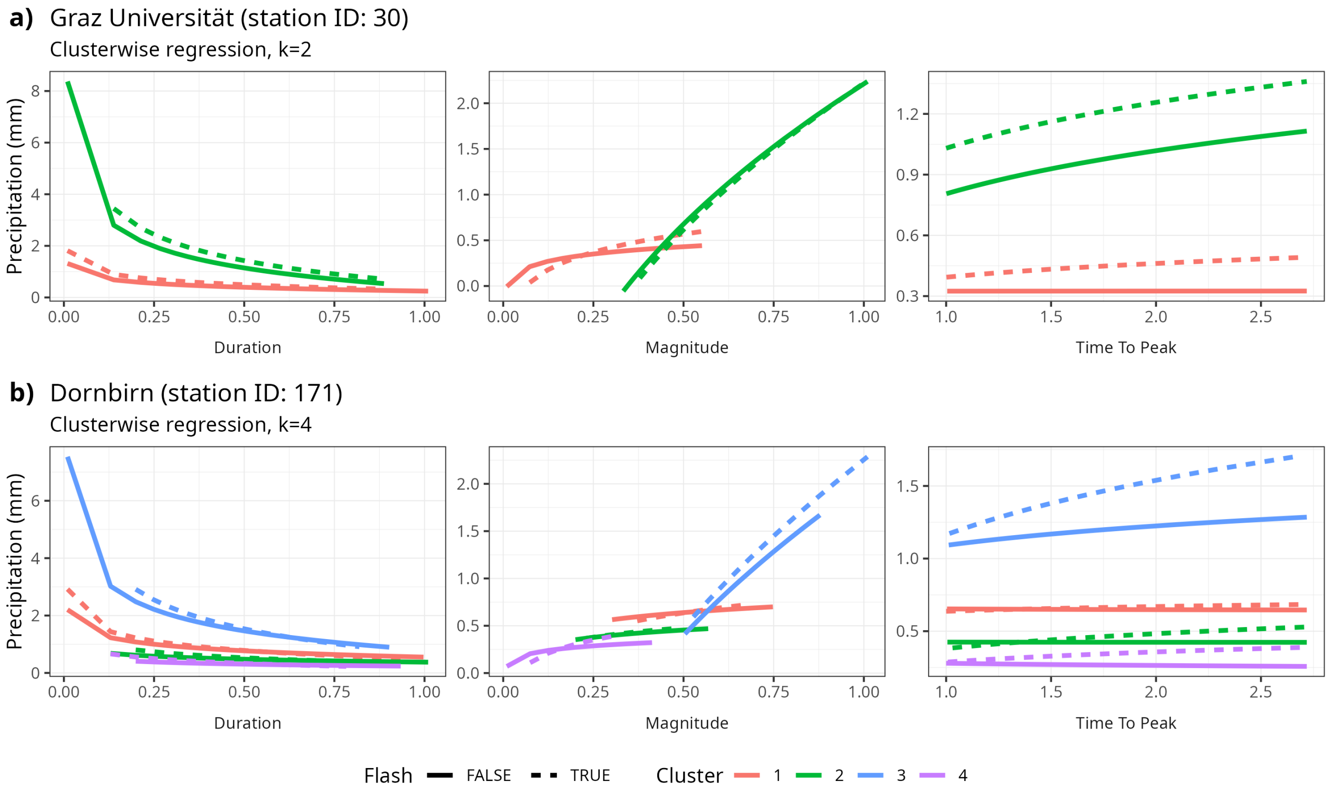

Figure 4. Partial Dependence Plots showing the marginal effect of event characteristics on predicted 10-minute precipitation in the clusterwise regression models. Solid lines denote events without lightning, dashed lines denote events with lightning. Colour indicates cluster membership. Rows: a) Graz Universität; b) Dornbirn.

3.2 Event characteristic effects within clusters

To examine how event characteristics influence the obtained clusters, we visualise their marginal effects using partial dependence plots (Friedman, 2001) (Fig. 4). The figure shows consistent patterns across both stations. Longer event `duration` is associated with lower predicted 10-min precipitation: in each station, one cluster shows a markedly steeper decline than the others (Dornbirn cluster 3: -37.9%; Graz cluster 2: -41.7%, across the observed range of `duration`), while the remaining clusters show comparatively flat relationships. (Dornbirn: -13.5% to -2.9%; Graz: -9.5%). Greater event `magnitude` is associated with higher predicted 10-min precipitation, again with one cluster per station showing a substantially steeper increase (Dornbirn cluster 3: +18.0%; Graz cluster 2: +23.1%) than the rest (Dornbirn: +1.1% to +2.5%; Graz: +6.0%).

The effect of the `time_to_peak` is more subtle, with weak slopes across most clusters (-0.2% to +1.7% change in predicted precipitation across the observed range) - markedly smaller than the `duration` and `magnitude` effects.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Table 1. Unconditional distribution-fit metrics (Kullback–Leibler divergence D_{KL}, root mean squared error RMSE and mean absolute error MAE of the CDF) for the single-Gamma baseline, partitioning clustering, and clusterwise regression. Values in parentheses indicate the percentage change relative to the single-Gamma baseline; negative values denote a reduction of the respective error metric, i.e. an improved fit.

Metric	Station	Single Gamma	Partitioning	Clusterwise regression
D_{KL}	Dornbirn	0.236	0.178 (-25%)	0.063 (-73%)
D_{KL}	Graz Universität	0.559	0.323 (-42%)	0.190 (-66%)
RMSE	Dornbirn	0.013	0.012 (-12%)	0.007 (-44%)
RMSE	Graz Universität	0.021	0.015 (-26%)	0.010 (-52%)
MAE	Dornbirn	0.004	0.003 (-19%)	0.001 (-65%)
MAE	Graz Universität	0.007	0.005 (-32%)	0.003 (-56%)

Events with lightning show consistently, though modestly, higher predicted precipitation than non-lightning events across all clusters and both stations (+0.04% to +1.8%). Lightning also modulates the `duration` relationship: in the cluster with the steepest decline at each station, lightning markedly dampens it (Dornbirn cluster 3: from -37.9% to -15.5%; Graz cluster 2: from -41.7% to -20.4%), whereas in the flatter clusters lightning instead slightly steepens the decline. For `time_to_peak`, lightning consistently strengthens the already weak positive slope across all clusters and both stations (by 0.1 to 3.0 percentage points), though those increases remain small in absolute terms. The effect of lightning on the `magnitude` relationship is smaller and inconsistent in direction across clusters and stations.

3.3 Improvement of theoretical distribution fit

Having discussed the characteristics of the obtained precipitation events, we now turn to the primary aim of this paper: whether the additional steps of clustering and clusterwise distribution aggregation sufficiently improve the theoretical Gamma distribution fit.

We first assess the improvements visually in Fig. 5. The figure again shows the results per station (rows) and clustering approach (columns). The ECDFs of the 10-min precipitation time series are shown in black, the fit of the single-Gamma distribution fitted to the complete time series in grey, and the clusterwise-aggregated Gamma distribution fit is shown in dark red. Overall, theoretical distribution fits are better for station Dornbirn than for Graz. Under the partitioning approach, the clusterwise-aggregated fit yields little to no improvement over the single-Gamma baseline. In contrast, under the clusterwise regression approach, the clusterwise-aggregate distribution shows substantial improvement across the entire precipitation range, including both the bulk of the distribution and its tails. This improvement is more pronounced in Graz than in Dornbirn, although in Dornbirn – where the baseline already performs well – the clusterwise-aggregated fit still offers visible gains.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

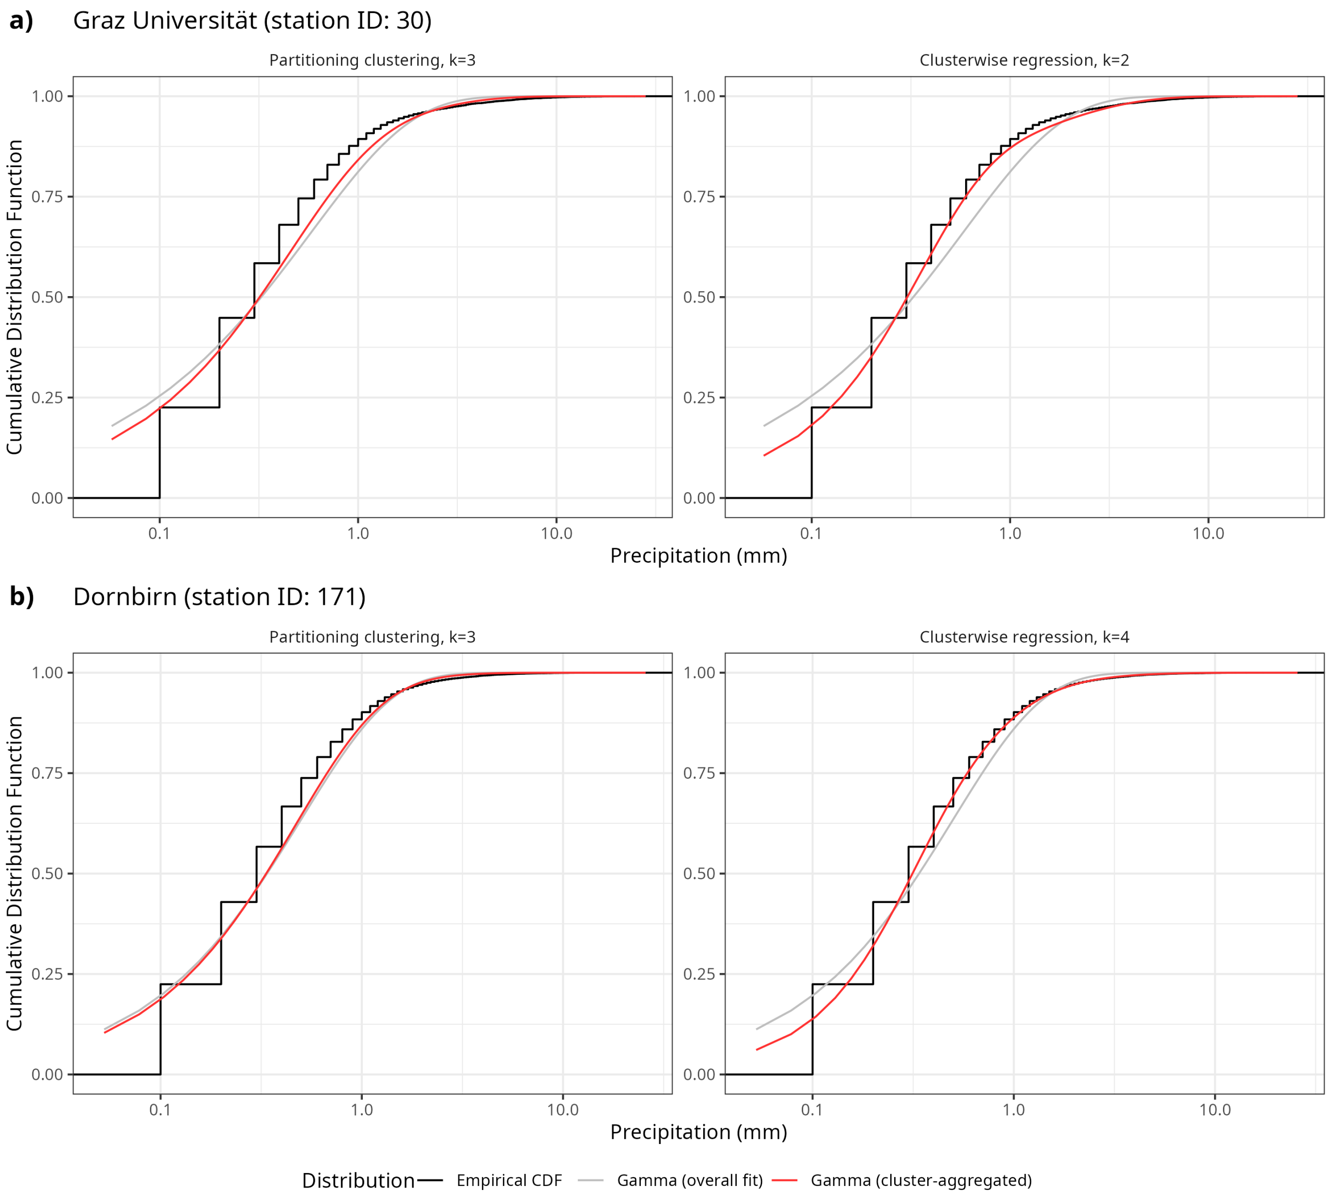

Figure 5. Visual comparison of theoretical Gamma distribution fit to precipitation ECDFs (black). Grey: single-Gamma fit to the complete precipitation time series. Dark red: cluster-aggregated Gamma obtained by weighting clusterwise fits. Panels are arranged by station (rows: top Graz Universität, bottom Dornbirn) and clustering approach (columns: left partitioning, right clusterwise regression).

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

We now quantify these improvements using three metrics: Kullback-Leibler divergence D_{KL}, the root mean squared error (RMSE) of the CDF, and the mean absolute error (MAE) of the CDF; all shown in Table 1. For all metrics, lower values indicate better fit. A complementary evaluation of the MAE restricted to empirical quantile bands is provided in Table C1 in Appendix C.

The D_{KL} results confirm the visual findings. The baseline D_{KL} is identical under both clustering approaches, as it reflects the single-Gamma fit to the complete dataset irrespective of clustering. For station Dornbirn, characterised by rainfall events with long duration and low intensity, the baseline fit is better, with a D_{KL} 58% lower than at Graz. The cluster-aggregated distributions improved upon the single-Gamma in all cases – both stations and both clustering approaches – with a mean D_{KL} reduction of 52%. The improvement is substantially larger for the clusterwise regression approach: a 66% reduction for station Graz and 73% for station Dornbirn (both relative to the baseline).

The RMSE and quantile-band MAE results are consistent with this picture: for both metrics and stations, the clusterwise regression's cluster-aggregated distribution outperforms the single-Gamma baseline. The band-restricted MAE (Table C1) confirms that the improvement spans the full distribution rather than being confined to a particular range. The one exception is the $[Q_{20}, Q_{40}]$ band at Dornbirn, where the gain under clusterwise regression is marginal (1%) and smaller than under the partitioning approach (7%).

To further illustrate where in the distribution these improvements occur, Fig. C1 in Appendix C shows P–P plots comparing the smoothed empirical CDF with the theoretical CDFs over the full probability range: while Fig. 5 shows the distributions themselves, the P–P plots focus on the residual deviations from the empirical CDF. The clusterwise regression's cluster-aggregated distribution tracks the diagonal closely at both stations, outperforming the single-Gamma at 96.8% of evaluated grid points for Graz and 98.7% for Dornbirn. The only contiguous region where the single-Gamma marginally outperforms is a narrow upper-tail band at station Graz (empirical CDF 95.2–96.7%, 27 of 1 000 grid points); no comparable contiguous region is observed at Dornbirn.

4 Discussion

This study asked whether the recurrent lack of fit of single-family rainfall distributions is primarily a symptom of process heterogeneity rather than an inherent shortcoming of the distributional families themselves. Our results indicate a tentative first "yes".

In a comparative case study of two meteorologically divergent Austrian weather stations (compare with Seibert et al., 2007), we introduced a two-step, process-aware workflow that first separates precipitation events into types and then combines cluster-specific subdistributions into a discrete mixture to represent the overall event-level distribution. The central contribution is methodological: an outcome-guided, clusterwise distributional regression that is supervised by the event-level precipitation distribution and uses event characteristics as covariates. This aligns with Gumbel's notion of homogeneity as "subject to a common set of forces" (Gumbel, 1941) and provides a natural statistical mechanism to approximate meteorological process types with distributional consequences.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

The comparison of clustering strategies illustrates this point. Partitioning clustering using k-means yields clusters that align closely with the axes of the event characteristics `duration`, `magnitude/severity`, `intensity`, `time-to-peak` and `lightning`, and are therefore readily interpretable. However, the cluster-aggregated Gamma derived from these groups delivers only modest improvements over the single-Gamma baseline. By contrast, clusterwise distributional regression produced event groups that are less tightly aligned with the observed characteristics alone, yet their weighted mixture consistently improves the theoretical fit across the entire probability range. The visual diagnostics (ECDFs and P–P plots; Figs. 3 and C1) and quantitative metrics (D_{KL}, CDF RMSE, and band-restricted MAE; Table 1) show that the gains extend to both the bulk and the tails, with larger improvements at Graz and still visible gains at Dornbirn where the baseline already performs comparatively well. Together, these results indicate that salient structure in event-level precipitation distributions is not fully spanned by the observed characteristics and that supervised clustering captures additional, outcome-relevant variation.

A common response to poor distributional fit in event-level models has been to replace the single-Gamma with heavier-tailed or more flexible families: Wilks (1998) and Li et al. (2013) advocated mixed Exponential distributions as an improvement over the standard Gamma, while Papalexiou and Koutsoyiannis (2012) and Papalexiou and Koutsoyiannis (2016) found that the Generalised Gamma (GG) and the Burr Type XII (BurrXII) achieve better coverage of the upper tail of wet-day precipitation. That line of work targets the "which family?" question on a pooled sample that implicitly mixes processes and assumes i.i.d. deviations. The present study pursues a complementary route: it first asks *which process* generated an event, separates events into more homogeneous types accordingly, and only then fits distributions within these types. Statistically, this is a classical finite-mixture exercise; physically, it acknowledges that different generating mechanisms (e.g., convective versus stratiform) imprint different distributions. The two routes are complementary rather than competing: our framework is distribution-agnostic and can embed any parametric family at the cluster level. Extending the clusterwise mixtures to include mixed Exponential, Generalised Gamma, or Burr XII within clusters is a natural next step, particularly to refine tail behaviour. An especially promising direction is to allow clusters to adopt different distributional families, for instance a lighter-tailed family for stratiform-like clusters and a heavier-tailed one for convective-like clusters – which could further improve the aggregated fit while providing new insight into the physical nature of distinct precipitation regimes.

The proposed event-type separation also sits naturally alongside existing hydrometeorological typologies that rely on extensive meteorological inputs and synoptic classifications, such as circulation or weather types and regional station groupings. Our approach uses information available at the event scale in the target series—duration, magnitude/severity, intensity, and timing—and augments it with lightning, yielding a simple, portable and effective event-type separation. The resulting groups align with hydrometeorological reasoning: at Dornbirn, separation tracks duration and magnitude/severity, consistent with more stratiform-like behaviour, while at Graz, intensity and lightning play a stronger role, consistent with a higher convective share. The resulting discrete mixture capitalises on this grouping to construct an overall distribution that better matches heterogeneous precipitation regimes.

Two limitations merit discussion. First, no general definition exists of what constitutes a precipitation event; see Dunkerley (2008) for a comparative review of event definitions via different minimum inter-event times (MIT) and other criteria, concluding that different MIT values lead to substantially different descriptive outcomes. We assessed robustness by repeating the

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

analyses for minimum inter-event times of 3, 6, 12, 18 and 24 h. Absolute fit (D_{KL}) generally degraded with inreasing MIT for both single-Gamma and cluster-aggregated distributions – as expected, since longer minimum inter-event times pool more rainfall into fewer, longer events and thereby increase within-event heterogeneity – yet the qualitative ordering remained stable: clusterwise regression outperformed partitioning clustering, and both outperformed the single Gamma, at both stations. This suggests that the main conclusions are robust to reasonable choices of MIT, while reinforcing the need to report event-definition choices transparently. Second, the FAMD and partial-dependence results indicate that unobserved factors influence event-level distributions beyond the characteristics considered here. Physically grounded covariates that are operationally accessible include measures of the convective environment (e.g., CAPE/CIN, precipitable water, low-level jet indices, wind shear), synoptic context (e.g., circulation or weather-type classifications from reanalyses, frontogenesis metrics), storm-structure proxies (e.g., radar convective/stratiform classification, echo-top height), orographic exposure, and more detailed lightning descriptors (e.g., flash density and timing within events). Incorporating such drivers could sharpen cluster definitions and further improve mixture fits, particularly in the tails.

We chose the single-Gamma as a transparent baseline to isolate the value added by clustering; as noted above, any family – or set of families – can be deployed at the cluster level, so the approach provides a vehicle for using flexible distributions where they are most appropriate rather than a single compromise over mixed processes. Scaling up the analysis to a broader set of stations and climatic regions will test portability and clarify how cluster composition varies with regime. Regional or hierarchical mixtures that share information across stations could stabilise parameter estimates for data-sparse clusters. With regard to generalisation, we mitigated overfitting by selecting the number of clusters k via bootstrap stability (ARI). For deployment, we recommend explicit cross-validation that preserves event blocks and station-level splits, and we advocate reporting out-of-sample diagnostics, including P–P plots and band-restricted errors.

The implications for applications are direct. Improved event-level distributions should benefit rainfall–runoff simulations, flood and erosion modelling, and the design of rainfall generators, particularly in regimes where process mixing is strong. In climate and weather generation, event-type mixtures provide a principled way to couple occurrence models with intensity models and to propagate regime-specific change signals in scenario analyses. More broadly, the results support a process-based view of distribution modelling in which classification matters: outcome-guided, clusterwise distributional regression provides a supervised, process-consistent way to separate events and markedly improves fit when combined into a discrete mixture. Addressing process heterogeneity via mixtures thus complements, rather than replaces, the move to more flexible single families, and both can be combined within one framework. Finally, event characteristics derived from the target series, augmented by lightning, already contain a wealth of process information and yield practically useful event types. A clear path forward is to add physically grounded covariates, allow different families across clusters, extend and regionalise with rigorous out-of-sample validation, and integrate the approach into rainfall generators and hydrological models where improved realism in the tails matters most.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

5 Conclusions

The present work demonstrates that event-level analysis of precipitation time series, and the derivation of process-aware, event-type mixture distributions substantially improves distributional fit compared to traditional single-distribution approaches. Across two Austrian stations (Graz Universität and Dornbirn, 10-min records, 2010–2022), clusterwise mixtures reduced D_{KL} by 66-73% relative to a single-Gamma baseline and improved ECDF and P–P diagnostics, with results robust to the choice of minimum inter-event time.

Our comparative analysis reveals that "outcome-guided" clustering via finite mixtures of Gamma regressions (clusterwise distributional regression) proved a particularly successful approach, delivering the best overall fit, but partitioning clustering remains a valuable tool for interpretable event-type delineation. While demonstrated here for two Austrian weather stations, and on a single distribution family, we believe that this methodology has considerable potential for broader application in hydrological statistics, particularly for applications requiring accurate representation of precipitation distribution tails, such as flood risk assessment and climate change impact studies. The framework is distribution-agnostic and portable (relying on event-scale information and lightning); extensions to heavier-tailed families (e.g. Generalised Gamma, Burr XII) and additional covariates are straightforward. Future work should assess multi-site transferability, incorporate hierarchical pooling for sparse sites, and evaluate non-stationarity under climate change to further strengthen tail inference for risk-critical applications.

Code and data availability. The precipitation time series are freely available at GeoSphere Austria, https://data.hub.geosphere.at/dataset/klima-v1-10min (DOI: 10.60669/8fya-7x87). Lightning data were provided by ALDIS (Austrian Lightning Detection and Information System); methodological details on detection efficiency and location accuracy are given in Schulz et al. (2005); Diendorfer et al. (1998). Analysis code and reproducible scripts are available at Github (https://github.com/boku-stat/ClusteringRainfall/tree/v1.0.0) and archived on Zenodo (https://doi.org/10.5281/zenodo.21358297).

Author contributions. L.O.M. led the formal analysis, developed the software, and coordinated the visualisation. N.B.Ö. performed the initial data processing and exploratory analyses and contributed to the exploitation of results. J.L. contributed to software implementation and visualisation. G.L. conceived and supervised the study in the role of senior author. All authors contributed to writing and reviewing the manuscript.

Competing interests. All authors declare that they have no conflicts of interest.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Acknowledgements. We thank the BOKU University's high-performance computing clusters for the computational resources provided, which greatly facilitated the simulations in this study. This research has been supported by the Austrian Climate and Energy Fund under the program "ACRP13" (grant no. KR20AC0K17974).

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

References

Beven, K. J.: Rainfall-runoff modelling: the primer, John Wiley & Sons, https://books.google.at/books?id=oB5G3ekgxpgC, 2012.

Diendorfer, G., Schulz, W., and Rakov, V. A.: Lightning characteristics based on data from the Austrian lightning locating system, IEEE Transactions on Electromagnetic Compatibility, 40, 452–464, https://doi.org/10.1109/15.736203, 1998.

Diendorfer, G. et al.: Lightning location systems (LLS), in: IX International Symposium on Lightning Protection, pp. 1–13, 2007.

Drost, H.: Philentropy: Information Theory and Distance Quantification with R, Journal of Open Source Software, 3, 765, https://doi.org/10.21105/joss.00765, 2018.

Dunkerley, D.: Identifying individual rain events from pluviograph records: a review with analysis of data from an Australian dryland site, Hydrological Processes, 22, 5024–5036, https://doi.org/10.1002/hyp.7122, 2008.

Elshamy, M. E., Wheater, H. S., Gedney, N., and Huntingford, C.: Evaluation of the rainfall component of a weather generator for climate impact studies, Journal of Hydrology, 326, 1–24, https://doi.org/10.1016/j.jhydrol.2005.09.017, 2006.

Friedman, J. H.: Greedy function approximation: a gradient boosting machine, Annals of statistics, pp. 1189–1232, https://doi.org/10.1214/aos/1013203451, 2001.

Grün, B. and Leisch, F.: Fitting Finite Mixtures of Generalized Linear Regressions in R, Computational Statistics & Data Analysis, 51, 5247–5252, https://doi.org/10.1016/j.csda.2006.08.014, 2007.

Grün, B. and Leisch, F.: flexmix: Flexible Mixture Modeling, https://doi.org/10.32614/CRAN.package.flexmix, r package version 2.3-20, 2025.

Gumbel, E. J.: The return period of flood flows, The annals of mathematical statistics, 12, 163–190, https://www.jstor.org/stable/2235766, 1941.

Hosking, J. R.: L-moments: analysis and estimation of distributions using linear combinations of order statistics, Journal of the Royal Statistical Society Series B: Statistical Methodology, 52, 105–124, https://doi.org/10.1111/j.2517-6161.1990.tb01775.x, 1990.

Hosking, J. R. M.: L-Moments, https://doi.org/10.32614/CRAN.package.lmom, r package, version 3.2, 2024.

Hubert, L. and Arabie, P.: Comparing partitions, Journal of classification, 2, 193–218, https://doi.org/10.1007/BF01908075, 1985.

Kaufman, L. and Rousseeuw, P. J.: Finding Groups in Data: An Introduction to Cluster Analysis, Wiley Series in Probability and Statistics, John Wiley & Sons, Inc., Hoboken, New Jersey, https://doi.org/10.1002/9780470316801, 1990.

Kullback, S. and Leibler, R. A.: On Information and Sufficiency, The Annals of Mathematical Statistics, 22, 79–86, http://www.jstor.org/stable/2236703, 1951.

Laaha, G., Gauster, T., Tallaksen, L. M., Vidal, J.-P., Stahl, K., Prudhomme, C., Heudorfer, B., Vlnas, R., Ionita, M., Van Lanen, H. A., et al.: The European 2015 drought from a hydrological perspective, Hydrology and Earth System Sciences, 21, 3001–3024, https://doi.org/10.5194/hess-21-3001-2017, 2017.

Laaha, G., Laimighofer, J., Özcelik, N. B., and Fischer, S.: Exploring Process Heterogeneity in Environmental Statistics: Examples and Methodological Advances, Austrian Journal of Statistics, 54, 124–149, https://doi.org/10.17713/ajs.v54i3.2101, 2025.

Lê, S., Josse, J., and Husson, F.: FactoMineR: A Package for Multivariate Analysis, Journal of Statistical Software, 25, 1–18, https://doi.org/10.18637/jss.v025.i01, 2008.

Leisch, F.: A toolbox for k-centroids cluster analysis, Computational statistics & data analysis, 51, 526–544, https://doi.org/10.1016/j.csda.2005.10.006, 2006.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Li, Z., Brissette, F., and Chen, J.: Finding the most appropriate precipitation probability distribution for stochastic weather generation and hydrological modelling in Nordic watersheds, Hydrological Processes, 27, 3718–3729, https://doi.org/10.1002/hyp.9499, 2013.

Papacharalampous, G., Dallan, E., Armon, M., Saha, J., Price, C., Borga, M., and Marra, F.: Precipitation-driven typology of storms in the Alps, https://doi.org/10.31223/X56M9W, 2025.

Papalexiou, S. and Koutsoyiannis, D.: Battle of extreme value distributions: A global survey on the extreme daily rainfall, Water Resources Research, 49, 187–201, https://doi.org/10.1029/2012WR012557, 2013.

Papalexiou, S., Koutsoyiannis, D., and Makropoulos, C.: How extreme is extreme? An assessment of daily rainfall distribution tails, Hydrology and Earth System Sciences, 17, 851–862, https://doi.org/10.5194/hess-17-851-2013, 2013.

Papalexiou, S. M. and Koutsoyiannis, D.: Entropy based derivation of probability distributions: A case study to daily rainfall, Advances in Water Resources, 45, 51–57, https://doi.org/10.1016/j.advwatres.2011.11.007, 2012.

Papalexiou, S. M. and Koutsoyiannis, D.: A global survey on the seasonal variation of the marginal distribution of daily precipitation, Advances in Water Resources, 94, 131–145, https://doi.org/10.1016/j.advwatres.2016.05.005, 2016.

Petroselli, A., Apollonio, C., De Luca, D. L., Salvaneschi, P., Pecci, M., Marras, T., and Schirone, B.: Comparative evaluation of the rainfall erosivity in the rieti province, central italy, using empirical formulas and a stochastic rainfall generator, Hydrology, 8, 171, https://doi.org/10.3390/hydrology8040171, 2021.

R Core Team: R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, https://www.R-project.org/, 2025.

Richardson, C. W. and Wright, D. A.: WGEN: A Model for Generating Daily Weather Variables, Tech. Rep. ARS-8, United States Department of Agriculture, Agricultural Research Service, 1984.

Schulz, W.: Location accuracy improvements of the Austrian lightning location system during the last 10 years, in: Proceedings of the 9th Asia-Pacific International Conference on Lightning (APL), Nagoya, Japan, pp. 23–27, 2015.

Schulz, W., Cummins, K., Diendorfer, G., and Dorninger, M.: Cloud-to-ground lightning in Austria: A 10-year study using data from a lightning location system, Journal of Geophysical Research: Atmospheres, 110, D09 101, https://doi.org/10.1029/2004JD005332, 2005.

Seibert, P., Frank, A., and Formayer, H.: Synoptic and regional patterns of heavy precipitation in Austria, Theoretical and Applied Climatology, 87, 139–153, https://doi.org/10.1007/s00704-006-0198-8, 2007.

Tallaksen, L. M. and Van Lanen, H. A.: Hydrological drought: processes and estimation methods for streamflow and groundwater, Elsevier, 2023.

Thom, H. C.: A note on the gamma distribution, Monthly weather review, 86, 117–122, https://doi.org/10.1175/1520-0493(1958)086<0117:ANOTGD>2.0.CO;2, 1958.

Todorovic, P. and Zelenhasic, E.: A Stochastic Model for Flood Analysis, Water Resources Research, 6, 1641–1648, https://doi.org/10.1029/WR006i006p01641, 1970.

Vásquez, C., Klik, A., Stumpp, C., Laaha, G., Strauss, P., Özcelik, N. B., Pistotnik, G., Yin, S., Dostal, T., Gaona, G., et al.: Rainfall erosivity across Austria's main agricultural areas: Identification of rainfall characteristics and spatiotemporal patterns, Journal of Hydrology: Regional Studies, 53, 101 770, https://doi.org/10.1016/j.ejrh.2024.101770, 2024.

Wickham, H.: ggplot2: Elegant Graphics for Data Analysis, Springer-Verlag New York, https://ggplot2.tidyverse.org, 2016.

Wilks, D. S.: Multisite generalization of a daily stochastic precipitation generation model, Journal of Hydrology, 210, 178–191, https://doi.org/10.1016/S0022-1694(98)00186-3, 1998.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Wilks, D. S. and Wilby, R. L.: The weather generation game: a review of stochastic weather models, Progress in physical geography, 23, 329–357, https://doi.org/10.1177/030913339902300302, 1999.

Yaqub, A., Seibert, P., and Formayer, H.: Diurnal precipitation cycle in Austria, Theoretical and Applied Climatology, 103, 109–118, https://doi.org/10.1007/s00704-010-0281-z, 2011.

Yevjevich, V. M. et al.: An objective approach to definitions and investigations of continental hydrologic droughts, vol. 23, Colorado State University Fort Collins, CO, USA, http://hdl.handle.net/10217/61303, 1967.

Zaghloul, M., Papalexiou, S. M., Elshorbagy, A., and Coulibaly, P.: Revisiting flood peak distributions: A pan-Canadian investigation, Advances in Water Resources, 145, 103 720, https://doi.org/10.1016/j.advwatres.2020.103720, 2020.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Appendix A: Summary statistics for event characteristics

Table A1. Summary statistics of the calculated event characteristics for the two stations. 171: Dornbirn, 30: Graz Universität.

Station	Characteristic	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
171	Severity [mm]	1.30	3.90	8.60	14.65	17.88	193.60
	Magnitude [mm]	0.30	0.80	1.40	2.13	2.60	25.70
	Duration [min]	10.0	110.0	280.0	408.7	550.0	3330.0
	Intensity [mm min^{-1}]	0.006591	0.022000	0.034286	0.048951	0.056360	0.470000
	Time To Peak []	0.007299	0.166667	0.400000	0.440164	0.677079	1.000000
	Flash []	FALSE:764	TRUE :194				
30	Severity [mm]	1.30	3.00	6.70	11.08	14.50	113.30
	Magnitude [mm]	0.300	0.800	1.400	2.724	3.400	28.100
	Duration [min]	10.0	60.0	170.0	267.1	360.0	2230.0
	Intensity [mm min^{-1}]	0.00641	0.02387	0.04107	0.06833	0.07714	0.66857
	Time To Peak []	0.008475	0.166667	0.409091	0.441883	0.666667	1.000000
	Flash []	FALSE:467	TRUE :238				

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

Appendix B: Results of the (global) stability analyses

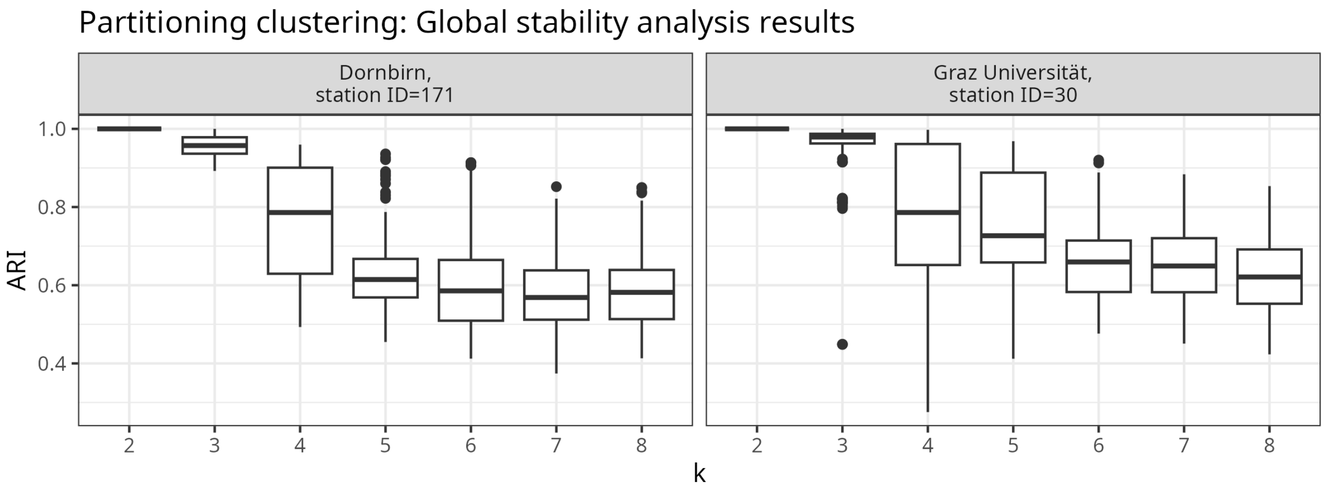

Figure B1. Global stability analysis results for partitioning clustering solutions in the two stations. x-axis: k, number of clusters tested. y-axis: Adjusted Rand Indices (ARIs) obtained by comparing 100×2 bootstrapped clustering runs.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

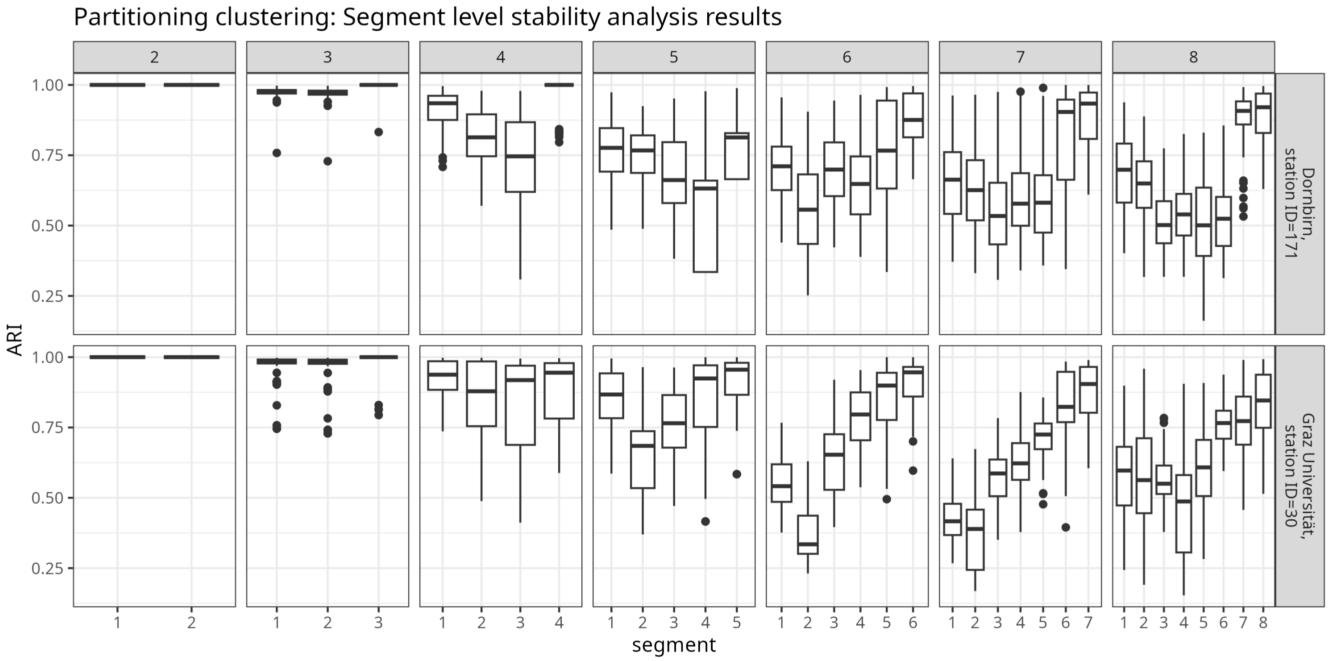

Figure B2. Segment level stability analysis results for partitioning clustering solutions in the two stations. Row labels: tested total number of clusters k, x-axis: each cluster segment of the tested solution. y-axis: Adjusted Rand Indices (ARIs) obtained by comparing 100×2 bootstrapped clustering runs, for each segment (e.g. how stable is the specific segment?).

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

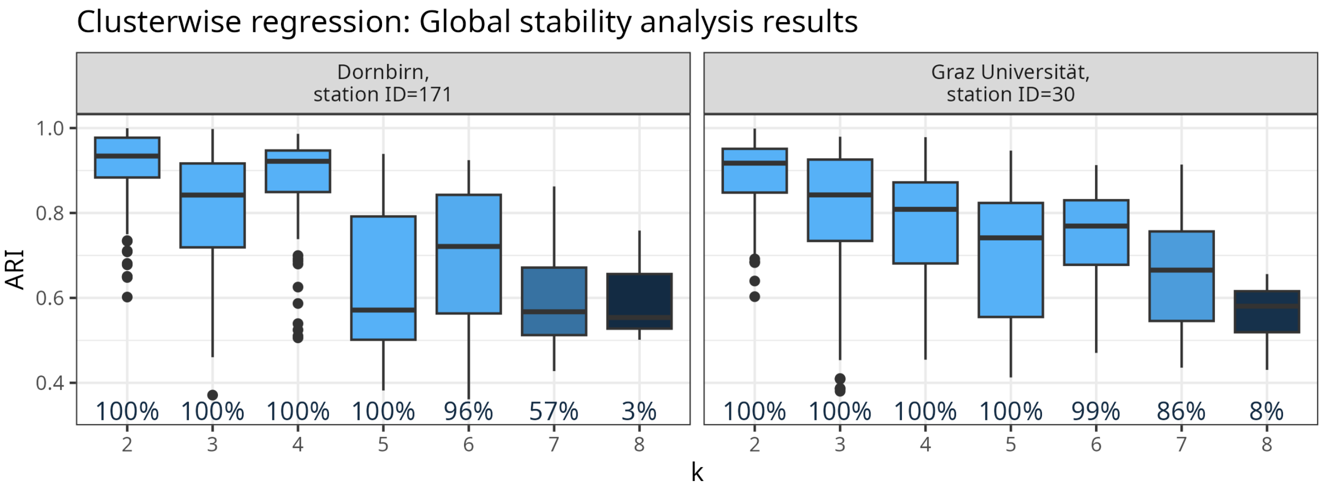

Figure B3. Global stability analysis results for clusters obtained via clusterwise regression in the two stations. x-axis: k, number of clusters tested. y-axis: Adjusted Rand Indices (ARIs) obtained by comparing 100×2 bootstrapped clustering runs. Color scale: Percentage (=count, as $nboot = 100$) of bootstrap runs that converged to the specified k (exact value is indicated in the text labels).

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

465 ## Appendix C: Comparing the fit of the theoretical (aggregated) Gamma distributions against the full distribution range via (E)CDFs and quantile bands

C1 P–P - plots of theoretical CDFs against ECDF

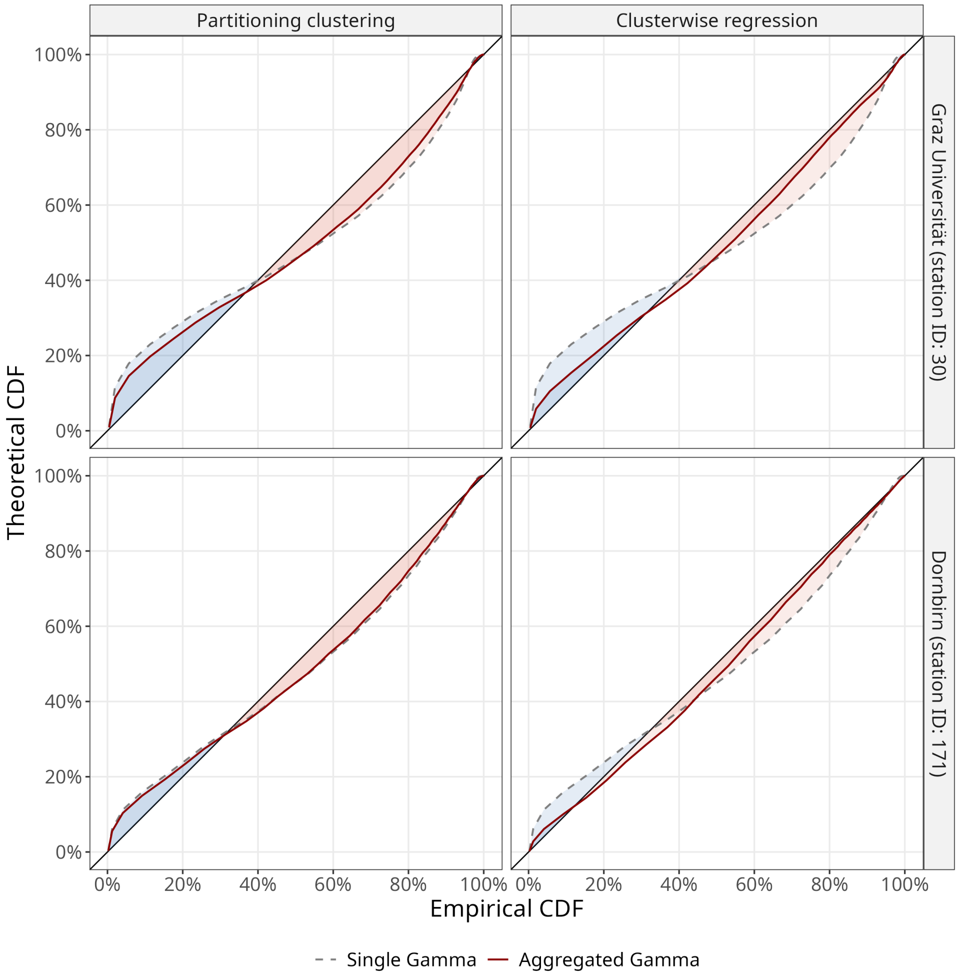

Figure C1. P–P plots comparing the smoothed empirical cumulative distribution function (CDF) with theoretical CDFs. Columns show the clustering approach (left partitioning, right clusterwise regression) and rows show the stations (top Graz Universität, bottom Dornbirn). The diagonal indicates a perfect fit. Blue shading marks regions where the theoretical CDF exceeds the empirical (overestimation of low values); red shading marks the opposite. The cluster-aggregated Gamma mixture (solid) closely tracks the diagonal across the full probability range, outperforming the single-Gamma at 96.8% of grid points for Graz and 98.7% for Dornbirn.

https://doi.org/10.5194/egusphere-2026-4282
Preprint. Discussion started: 24 July 2026

C2 Table of band-restricted mean absolute CDF error MAEs

Table C1. Band-restricted mean absolute CDF error $\mathrm{MAE}_{[Q_a,Q_b]}$ by station for the single-Gamma baseline, partitioning clustering, and clusterwise regression, given in units of 10^{-3}. Values in parentheses indicate the percentage change relative to the single-Gamma baseline; negative values denote a reduction, i.e. an improved fit. $\mathrm{MAE}_{[Q_a,Q_b]}$ denotes the mean absolute CDF error restricted to the $[a,b]$th empirical percentile band. The $[Q_0,Q_{20}]$ band is omitted as it contains too few grid points for a meaningful evaluation.

Metric	Station	Single Gamma	Partitioning	Clusterwise regression
$\mathrm{MAE}_{[Q_{20},Q_{40}]}$	Dornbirn	47.9	44.7 (-7%)	47.3 (-1%)
$\mathrm{MAE}_{[Q_{20},Q_{40}]}$	Graz Universität	103.5	81.0 (-22%)	62.2 (-40%)
$\mathrm{MAE}_{[Q_{40},Q_{60}]}$	Dornbirn	47.7	46.9 (-2%)	41.2 (-14%)
$\mathrm{MAE}_{[Q_{40},Q_{60}]}$	Graz Universität	45.4	42.3 (-7%)	34.3 (-24%)
$\mathrm{MAE}_{[Q_{60},Q_{80}]}$	Dornbirn	65.4	55.1 (-16%)	17.4 (-73%)
$\mathrm{MAE}_{[Q_{60},Q_{80}]}$	Graz Universität	97.5	71.3 (-27%)	23.3 (-76%)
$\mathrm{MAE}_{[Q_{80},Q_{100}]}$	Dornbirn	2.3	1.7 (-26%)	0.5 (-78%)
$\mathrm{MAE}_{[Q_{80},Q_{100}]}$	Graz Universität	5.1	3.3 (-35%)	2.2 (-57%)