



Capturing the extremes: a quasi-comonotonicity-based algorithm for disaggregating daily to hourly rainfall

Carlos Correa¹, Alfonso Hernanz¹, Iván San-Felipe¹, Esteban Rodríguez-Guisado¹

¹Spanish Meteorological Agency (AEMET), Madrid, 28040, Spain

5 Correspondence to: Carlos Correa (ccorreag@aemet.es)

Abstract. Disaggregating daily precipitation data into hourly time scale is crucial for hydrological modelling, urban drainage design, and extreme rainfall risk assessment. This study presents Q-CODA, a novel Quasi-Comonotonicity-based Disaggregation Algorithm that leverages the quasi-comonotonic relationship between daily precipitation totals and their sub-daily maxima to generate hourly rainfall sequences consistent with observed extremes. The method first assumes the

10 Fréchet–Hoeffding upper bound copula to simulate sub-daily maximum precipitation values conditioned on daily totals, which serve as target constraints. An initial hourly pattern is obtained via a K-Nearest Neighbours (KNN) approach and subsequently refined through an iterative adjustment algorithm to ensure coherence with both the daily precipitation total and previously calculated multiple sub-daily target constraints. We evaluated Q-CODA through a rigorous 5-fold cross-validation over 91 meteorological stations across Spain, spanning 1996–2024. Performance was benchmarked against state-

15 of-the-art approaches including nearest-neighbour resampling methods, Poisson cluster-based rainfall generators, multiplicative random cascades models and deep learning techniques. Evaluation metrics were tailored to different aspects of the rainfall data: distributional distance (1-D Wasserstein), accuracy measures (MAE, Nash–Sutcliffe efficiency), and extreme quantiles (90th to 99.9th percentiles) were computed on the hourly maximum precipitation over 1-hour windows, focusing on the representation of rainfall extremes. In contrast, autocorrelation and rainfall event duration statistics were

20 calculated on the entire hourly rainfall time series to assess temporal coherence and event structure. Additionally, intensity-duration-frequency (IDF) curves were analysed. Results demonstrate that Q-CODA substantially improves the representation of rainfall extremes while maintaining temporal coherence and event structure. This approach offers a robust, data-driven framework for accurate sub-daily rainfall disaggregation, with significant implications for hydrometeorological applications and infrastructure design.

25 1 Introduction

Rainfall is one of the most variable components of the hydrological cycle, exhibiting pronounced fluctuations across spatial and temporal scales. While sub-hourly or hourly rainfall data are essential for simulating surface runoff, designing urban drainage infrastructure, and assessing flood risks (Schilling, 1991; Fowler et al., 2021; Haslinger et al., 2025), such high-resolution records are often unavailable or incomplete due to technical and logistical constraints. In contrast, daily



30 precipitation measurements are widely available from global meteorological networks, leading to a growing demand for
accurate temporal disaggregation methods that can transform daily totals into realistic sub-daily sequences (Takhellambam et
al., 2022; Ebers et al., 2024).

Temporal disaggregation, also known as temporal downscaling, refers to the process of generating high-resolution time
35 series from lower-resolution observations while preserving key statistical and hydrological characteristics. Traditional
approaches to precipitation disaggregation can be broadly categorized into two families: stochastic rainfall generators and
resampling-based methods. Stochastic generators such as Poisson cluster processes (Rodriguez-Iturbe et al., 1987; Onof et
al., 2000; Qin and Dai, 2024) or Markov chain-based schemes (Stoner and Economou, 2020; Vorobevskii et al., 2024)
simulate rainfall sequences by modelling the occurrence and intensity of rainfall events using probabilistic laws. While
40 theoretically appealing, these models often require extensive calibration and may underperform in replicating extremes
(Vorobevskii et al., 2024). Resampling-based methods such as K-Nearest Neighbors (KNN) and Methods of Fragments
(MoF) disaggregation rely on analogues from historical high-resolution observations (Pui et al., 2012; Alzahrani et al.,
2023). These approaches are nonparametric, easy to implement, and capable of preserving observed patterns; however, KNN
and MoF methods can suffer from limitations including sensitivity to the choice of neighbours and potential mismatch lag-1
45 autocorrelation (Li et al., 2018).

To address some of these challenges and improve the realism of disaggregation, newer approaches such as the
Microcanonical Multiplicative Random Cascade (MMRC) methodology and deep learning techniques have been
increasingly explored. MMRC models disaggregate rainfall by recursively partitioning precipitation amounts into smaller
50 time intervals using multiplicative random weights, preserving the overall mass while reproducing the multifractal and
scaling properties of rainfall (Förster et al., 2016; Müller and Haberlandt, 2018; Müller-Thomy, 2020; Maloku et al., 2023).
In parallel, deep learning methods have gained traction in rainfall disaggregation due to their ability to capture complex
nonlinear dependencies (Scher and Peßenteiner, 2021; Bhattacharyya and Saha, 2023; Oates et al., 2025). Hybrid models
combining Artificial Neural Networks (ANN) with clustering algorithms (e.g., ANN-K) have shown promise in improving
55 event-based performance and capturing nonlinear patterns. A recent study (Bhattacharyya et al., 2024) indicates that ANN-K
produces better results in conserving extreme rainfall than MMRC and Neyman-Scott Rectangular Pulse (NSRP) processes.
Nonetheless, such data-driven deep learning methods require large training datasets, risk overfitting, and often lack physical
interpretability, making their generalization across climatic regions challenging.

60 A particularly persistent difficulty in temporal disaggregation is the adequate reproduction of extreme rainfall intensities and
their temporal structure. Most existing methods either under-represent short-duration maxima or distort their alignment with
daily totals, which can lead to substantial biases in hydrological simulations and impact assessments. This motivates the need



for innovative strategies that explicitly account for the statistical dependence between daily and sub-daily precipitation features.

65

Copula theory (Sklar, 1996) provides a robust mathematical framework for modelling such dependencies, allowing the construction of joint distributions that preserve individual marginals while flexibly capturing inter-variable dependence. Copula-based models have been successfully applied to hydrological problems such as rainfall-runoff modelling, multivariate drought analysis, and spatial rainfall simulation (Naderi et al., 2022; Moradzadeh Rahmatabadi et al., 2023).

70 However, previous copula-based approaches to rainfall modelling have primarily focused on capturing the statistical dependence between precipitation intensities at different temporal or spatial scales, or on modelling joint distributions of variables such as storm duration and depth (Vorobevskii et al., 2024; Biswas and Saha, 2025). These approaches have not explicitly leveraged the quasi-comonotonic relationship between sub-daily precipitation maxima (e.g., 1-hour, 2-hour intensities) and the corresponding daily total as a structural constraint in the simulation of high-resolution rainfall time series.
75 This quasi-comonotonicity implies that extreme hourly accumulations tend to co-occur with large daily totals, imposing an upper limit on feasible sub-daily intensities.

In this study, we propose a novel assumption of the Fréchet–Hoeffding upper bound copula (Joe, 2005; Abdellatif et al., 2024), also known as comonotonicity copula, to simulate sub-daily maxima conditioned on daily totals, thereby generating
80 physically plausible upper envelopes that act as dynamic, event-specific constraints. These constraints complement the daily total in guiding the subsequent disaggregation process and allow for the generation of hourly rainfall sequences that better preserve the magnitude and timing of extreme events. This methodological innovation constitutes the core contribution of our disaggregation framework and distinguishes it from existing statistical or deep learning-based approaches. In this context, we present Q-CODA (Quasi-Comonotonicity-based Disaggregation Algorithm), a novel precipitation
85 disaggregation framework that leverages the quasi-comonotonic dependence between daily totals and sub-daily maxima to constrain and guide the generation of hourly rainfall series. The proposed method comprises three key components: (i) simulation of sub-daily maximum intensities assuming the Fréchet–Hoeffding upper bound copula conditioned on daily rainfall, (ii) seeding of initial hourly patterns using a K-Nearest Neighbours approach, and (iii) iterative adjustment to enforce consistency with daily totals and the simulated sub-daily maxima across multiple accumulation windows (1, 2, 6,
90 and 12 hours).

We evaluate Q-CODA on a high-resolution dataset from 91 meteorological stations distributed across diverse climatic regions of Spain over the period 1996–2024. The performance is benchmarked against state-of-the-art disaggregation methods, including KNN, ANN-K, MMRC, and Poisson-cluster rainfall generators. Evaluation metrics encompass
95 distributional distances (e.g., 1-D Wasserstein distance), error measures (e.g., MAE and Nash–Sutcliffe efficiency), as well as indicators focused on extremes (quantiles from the 90th to 99.9th percentile), temporal persistence (autocorrelation), and



rainfall event duration statistics. We also assess the capability of Q-CODA to reproduce intensity-duration-frequency (IDF) curves.

- 100 By integrating theoretical copula constraints with data-driven refinement, Q-CODA offers a hybrid approach that combines interpretability, realism, and robustness. The method is particularly suitable for applications where the preservation of extreme sub-daily rainfall events is critical and provides a valuable contribution to the growing toolkit of high-resolution precipitation generation techniques.
- 105 The remainder of this paper is organised as follows. Section 2 describes the dataset and study domain. Section 3 introduces the proposed Q-CODA framework, including the derivation of sub-daily constraints, the comonotonic transformation, and the iterative adjustment procedure. Section 4 presents the benchmark disaggregation methods used for comparison. Section 5 outlines the evaluation strategy, results of the evaluation are presented in Sect. 6, and finally main conclusions are summarized in Sect. 7.

110 2 Data and Study Area

- This study leverages a high-quality, high-resolution precipitation dataset from 91 meteorological stations maintained by the Spanish State Meteorological Agency (AEMET). The selected stations are distributed across the Iberian Peninsula and the Balearic Islands, ensuring extensive spatial coverage across diverse climatic zones of Spain (see Fig. 1). The dataset spans a 29-year period (1996–2024), with hourly precipitation records available at 0.1 mm resolution. Only stations with at least
- 115 90% data completeness were retained to ensure data robustness and minimize biases introduced by missing values. The geographic diversity of the station network covers coastal, mountainous, and inland regions, as well as areas subject to Mediterranean, Atlantic, and continental climatic influences. This variability in hydroclimatic regimes, ranging from convective summer storms to winter frontal systems, provides a rigorous testbed for evaluating the performance and generalizability of the proposed Q-CODA disaggregation framework across multiple rainfall-generating mechanisms.
- 120 Importantly, the scale of this study sets it apart from many recent rainfall disaggregation studies that rely on significantly smaller station networks, often basing methodological conclusions on fewer than 10 sites (e.g., Lee et al., 2022; Bhattacharyya et al., 2024; Chowdhury et al., 2025). Such limited spatial sampling can hinder the generalization of results, particularly in regions with high climate variability. By contrast, our analysis aligns with recent best practices emphasizing broad spatial validation and data diversity (e.g., Müller-Thomy, 2020; Ebers et al., 2024), which are crucial for evaluating
- 125 the robustness of disaggregation methods across different hydrometeorological contexts.

For visualization, Fig. 1 shows the geographic location of the 91 stations used in the analysis, highlighting the spatial representativeness of the dataset. All data were subjected to a thorough quality control process and days with clearly



erroneous or missing hourly values were excluded from the analysis. The full hourly time series were used to derive both
 130 daily precipitation totals (for disaggregation input) and sub-daily maxima over selected accumulation durations (1 h, 2 h, 6 h,
 and 12 h), which served as internal constraints and evaluation targets throughout the disaggregation framework.




Figure 1: Map showing the location of AEMET stations with available hourly precipitation records used in this study (elevation data: NOAA 2022; <https://doi.org/10.25921/fd45-gt74>).

3 Methodology

3.1 Conceptual Framework

Temporal disaggregation of daily precipitation into sub-daily time scales relies fundamentally on capturing the strong
 dependence between daily precipitation totals and sub-daily maxima. Empirical evidence shows a quasi-comonotonic
 relationship between these variables, meaning that high sub-daily maxima tend to coincide with large daily totals. This near-
 140 perfect positive dependence provides a natural constraint for disaggregation, ensuring physically plausible rainfall
 sequences.

Copulas are functions that couple multivariate distribution functions to their univariate marginals, enabling flexible
 modelling of dependence structures independently of marginal behaviour. Formally, for random variables X and Y with



145 marginal cumulative distribution functions (CDFs) $F_X(x)$ and $F_Y(y)$, Sklar's theorem (Sklar, 1996; Joe, 2005; Abdellatif et al., 2024) states that there exists a copula C such that the joint CDF $F_{X,Y}(x, y)$ can be expressed as:

$$F_{X,Y}(x, y) = C(F_X(x), F_Y(y)) , \quad (1)$$

The copula $C : [0, 1]^2 \rightarrow [0, 1]$ encapsulates the dependence structure between X and Y .

150 Among the family of copulas, the Fréchet–Hoeffding bounds define the theoretical limits of dependence. The upper bound copula C^+ , representing perfect positive dependence (comonotonicity), is defined as:

$$C^+(u, v) = \min(u, v) , \quad (2)$$

for $u, v \in [0, 1]$.

This copula implies that X and Y increase together almost surely, which aligns with the quasi-comonotonic behaviour
 155 observed between daily precipitation totals and sub-daily maxima. Assuming the Fréchet–Hoeffding upper bound copula allows us to generate the upper envelope or conditional upper bound of sub-daily maxima given daily totals. This approach provides a dynamic constraint that complements the daily precipitation total by restricting the range of possible sub-daily intensities to those consistent with observed extreme behaviour. Unlike previous rainfall disaggregation approaches that often treat daily totals and sub-daily maxima independently or model their dependence without explicit constraints, our
 160 method leverages this quasi-comonotonic relationship as a key structural feature. This enhances the physical realism of simulated hourly sequences, especially for extreme events, by ensuring that sub-daily maxima do not exceed plausible bounds conditional on the daily total. To illustrate this quasi-comonotonic dependence, we include Fig. 2 showing the empirical Spearman rank correlation coefficient (ρ_s) between sub-daily precipitation maxima aggregated over different time windows ($P_t^{\max}$) and the corresponding daily total precipitation (P_d).

165

Fig. 2 clearly shows empirical values of Spearman correlation approaching 1, confirming the near-perfect positive dependence, though decreases slightly for shorter durations, reflecting the increasing variability of shorter accumulation periods. This figure visually supports the theoretical basis of our methodology by quantifying the strength of the dependence exploited through the Fréchet–Hoeffding upper bound copula.

170




Figure 2: Spearman rank correlation coefficient (ρ_s) between sub-daily precipitation maxima aggregated over varying time windows ($P_t^{\max}$) and the corresponding daily total precipitation (P_d). Each boxplot summarizes the distribution across the 91 stations.

3.2 Sub-daily Maximum Simulation

Following the conceptual framework introduced in Sect. 3.1, the Fréchet–Hoeffding upper bound copula is accepted to simulate physically consistent sub-daily maximum precipitation intensities conditioned on daily totals. This step provides a set of dynamic constraints that ensure physically plausible short-duration rainfall intensities during the disaggregation process.

Let P_d denote the daily precipitation total and $P_\tau^{\max}$ the maximum accumulated precipitation over a sub-daily windows of duration $\tau \in \{1 \text{ h}, 2 \text{ h}, 6 \text{ h}, 12 \text{ h}\}$. Assuming quasi-comonotonicity, the joint distribution of $(P_d, P_\tau^{\max})$ is modelled using the Fréchet–Hoeffding upper bound copula C^+ , defined as:

$$F_{P_d, P_\tau^{\max}}(x, y) = \min(F_{P_d}(x), F_{P_\tau^{\max}}(y)), \quad (3)$$

Given a daily total $P_d = x$, we aim to simulate the corresponding sub-daily maxima $P_\tau^{\max}$ such that the pair $(x, P_\tau^{\max})$ lies on or below the upper copula bound. Since C^+ corresponds to perfect rank correlation, the conditional distribution of $P_\tau^{\max} | P_d = x$ can be approximated via comonotonic transformation, Eq. (4):

$$P_\tau^{\max} = F_{P_\tau^{\max}}^{-1}(F_{P_d}(x)), \quad (4)$$

where F_{P_d} and $F_{P_\tau^{\max}}$ are the empirical cumulative distribution functions (ECDFs) estimated from historical in each climatological season (DJF, MAM, JJA, or SON). This formulation enforces a strong dependence structure between daily totals and sub-daily peaks, while preserving the observed marginal distributions. For each simulated day, a vector of



conditional maxima $\mathbf{M} = (P_{1h}^{\max}, P_{2h}^{\max}, P_{6h}^{\max}, P_{12h}^{\max})$ is generated, where $P_{\tau}^{\max}$ denotes the maximum allowable precipitation over a sliding window of duration τ , consistent with the observed comonotonic relationship. These constraints are subsequently enforced in the disaggregation process (see Sect. 3.4), ensuring that simulated hourly series are consistent with both the daily total and realistic sub-daily intensities. This step is essential to preserve the magnitude and timing of extreme sub-daily events, which are critical for hydrological applications such as flood modelling, urban drainage, and climate impact assessment.

To evaluate the suitability of assuming the upper copula C^+ to estimate $\mathbf{M}$ via comonotonic transformation, defined in Eq. (4), we compared its performance against Isotonic Regression (IR), a standard non-parametric machine learning method that fits monotonic functions without assuming a specific functional form (De Leeuw, 1977; Chakravarti, 1989; Delong and Wüthrich, 2024). Its use as a benchmark is particularly relevant here, as both IR and the Comonotonic Transformation (CT) rely on the monotonicity assumption, albeit from different theoretical foundations: IR enforces monotonicity through piecewise-constant regression fits, while CT exploits comonotonicity explicitly via copula theory. Both approaches were trained and tested using a consistent 5-fold cross-validation scheme, applied independently at each station. To ensure reproducibility, we used the same non-overlapping temporal folds defined in Sect. 5, where hourly disaggregation performance is evaluated in detail. Figure 3 presents the bias (%) in estimating the 99th percentile of observed sub-daily maxima for each duration in the constraint vector $\mathbf{M}$, aggregated across 91 stations.




Figure 3: Bias (%) in estimating the 99th percentile of observed sub-daily precipitation maxima for each duration included in the constraint vector $\mathbf{M}$, aggregated across 91 stations.

While IR can adapt to nonlinearities in the data, it struggles to accurately capture the tail behaviour of precipitation distributions. The CT consistently achieves lower biases compared to IR, with the differences being more pronounced at shorter aggregation durations, indicating a closer match to the empirical distributions of extreme short-duration rainfall.



These distributional advantages make CT a more appropriate choice for integrating sub-daily intensity information into the
 220 disaggregation framework, as detailed in the following sections.

3.3 Initial Pattern Seeding via KNN

To generate an initial realistic sub-daily precipitation pattern consistent with a given daily total (P_d), we employ a K-Nearest
 225 Neighbours (KNN) approach. For each target day in the test set, we identify historically observed daily events from the
 training set with similar P_d magnitudes, constrained to the same climatological season (DJF, MAM, JJA, or SON).

Given a target daily total P_d^{target} , the Euclidean distance is computed between P_d^{target} and each training daily total P_d^{train}
 within the same season. The k closest neighbours are identified, and one neighbour is randomly sampled. The associated 24-
 hour precipitation pattern $\mathbf{h} = (h_1, h_2, \dots, h_{24})$ of the selected neighbour is scaled proportionally to match the target total:

$$230 \quad \mathbf{h} \leftarrow \mathbf{h} \times \frac{P_d^{\text{target}}}{\sum_{i=1}^{24} h_i}, \quad (5)$$

This scaled pattern constitutes the initial seeding for the iterative refinement described in Sect. 3.4, which adjusts sub-daily
 intensities to meet prescribed constraints on maxima over multiple accumulation intervals.

3.4 Iterative Adjustment Procedure

Starting from the obtained initial hourly precipitation vector $\mathbf{h}$ via KNN and target constraints composed of the daily total
 235 precipitation P_d and the obtained sub-daily maxima vector $\mathbf{M}$ via CT, the iterative adjustment procedure modifies $\mathbf{h}$ to satisfy
 these constraints simultaneously. At each iteration, for each aggregation duration $\tau \in \{1 \text{ h}, 2 \text{ h}, 6 \text{ h}, 12 \text{ h}\}$, the procedure
 identifies the contiguous time window $I_\tau = \{i, \dots, i + \tau - 1\}$, where i is the starting hour of that window (ranging from 1 to 25
 $-\tau$), for which the current sum is maximal.

$$S_\tau^{\text{max}} = \max_{i=1, \dots, 25-\tau} \sum_{j \in I_\tau} h_j, \quad (6)$$

240 The adjustment magnitude is defined as the difference between the corresponding element in $\mathbf{M}$ and the maximum sum in $\mathbf{h}$
 for duration τ :

$$\Delta_\tau = P_\tau^{\text{max}} - S_\tau^{\text{max}}, \quad (7)$$

If $|\Delta_\tau|$ exceeds a small threshold ε (0.04 mm), the values of $\mathbf{h}$ within the window I_τ are uniformly adjusted to reduce this
 difference:

$$245 \quad h_j \leftarrow \max\left(0, h_j + \frac{\Delta_\tau}{\tau}\right) \quad \forall j \in I_\tau, \quad (8)$$



This ensures that the sub-daily maxima are adjusted upwards or downwards, depending on whether the current maximum is below or above the target. This process is repeated until convergence criteria are met or a maximum number of iterations is reached. After adjusting all durations, the entire hourly vector $\mathbf{h}$ is scaled to satisfy the daily total constraint using Eq. (5) again. To enhance the temporal realism of the resulting precipitation pattern, a refined adjustment step is applied. This step improves the lag-1 autocorrelation of $\mathbf{h}$, a key indicator of temporal smoothness and persistence, through two complementary strategies. First, small precipitation values immediately preceding rainfall events are shifted forward to the subsequent hour when doing so increases autocorrelation without violating the constraints on sub-daily maxima or total daily precipitation. Second, within sliding windows of 3 to 5 hours, permutations of nonzero precipitation values are explored to find reorderings that further enhance autocorrelation while preserving the constraints. These refinements mitigate abrupt changes and isolated low values, yielding hourly sequences that better reflect realistic precipitation dynamics.

To handle extremes exceeding the training range (i.e., daily totals surpassing the historical maximum), Q-CODA employs only the KNN seed directly, omitting both the comonotonic transformation described in Sect. 3.2 and the iterative adjustment procedure described in this section. The overall methodology is summarized in the flowchart presented in Fig. 4.

Q-CODA




Figure 4: Flowchart summarizing the overall Q-CODA methodology. The diagram outlines the sequential steps involved in the disaggregation approach.



4 Benchmark Methods

To evaluate the performance of the proposed Q-CODA method, we compare it against established rainfall disaggregation techniques: KNN, ANN-K, Poisson-Cluster Disaggregation, and MMRC. These benchmark methods span both data-driven and stochastic approaches, reflecting the diversity of strategies commonly employed for sub-daily precipitation generation.

4.1 K-Nearest Neighbors (KNN)

The KNN method serves as an analog-based disaggregation approach. For each target daily precipitation value, historical days with similar daily totals are identified within the same climatological season (DJF, MAM, JJA, or SON), ensuring that the selected analogues reflect seasonal variability in rainfall dynamics. Similarity is assessed using Euclidean distance, and among the k closest matches, one historical 24-hour pattern is randomly selected. This pattern is then linearly scaled to match the target daily total. No additional sub-daily structure adjustment is performed. This method represents a baseline based on seasonal similarity and direct pattern reuse, as described in (Alzahrani et al., 2023), whose study shows that KNN ($k=30$) outperforms their multiplicative random cascade model and Hurst–Kolmogorov process-based model on a set of specific performance metrics and four locations.

4.2 Deep learning disaggregation model (ANN-K)

This method builds on the deep learning-based disaggregation framework proposed by (Bhattacharyya et al., 2024), combining unsupervised clustering via K-Means with a supervised Artificial Neural Network (ANN). The key motivation for this hybrid approach is to improve model skill in differentiating between low-intensity and high-intensity rainfall events, which are often poorly represented by traditional neural networks due to their tendency to regress towards the mean or zero-inflated values.

Preprocessing is performed using K-Means clustering on triplets of consecutive daily rainfall values: $(P_d^{(t-1)}, P_d^{(t)}, P_d^{(t+1)})$. The clustering is done in three-dimensional space to capture the temporal context around each day. The optimal number of clusters is selected using the Silhouette Coefficient method (Rousseeuw, 1987), ranging from $K = 2$ to 9. This categorical information (cluster labels) is used as an additional input to the ANN, allowing the model to better differentiate convective events, dry days, and transitional regimes. Our model introduces an architecture comprising a 4-layer feedforward neural network, defined as follows: (i) Input layer – a concatenated vector consisting of daily rainfall and the corresponding cluster label; (ii) First hidden layer – 64 neurons with ELU (Exponential Linear Unit) activation; (iii) Second hidden layer – 32 neurons with ELU activation; and (iv) Output layer – 24 neurons (representing hourly rainfall estimates), with a linear activation function. This structure enables the model to learn abstract representations of intra-daily rainfall patterns. The use of ELU (Exponential Linear Unit) activation functions helps to reduce vanishing gradients and maintain responsiveness for negative inputs, improving convergence during training. The ELU activation function was chosen by (Bhattacharyya et al.,



2024) over ReLU or LeakyReLU due to its improved numerical stability and convergence when dealing with sparse and
 295 skewed precipitation data. In addition, the formulation of the loss function has a decisive influence (Lopez-Gomez et al.,
 2023; Oates et al., 2025) on the model's skill to disaggregate precipitation events. To guide the ANN toward physically
 meaningful disaggregation, we define a composite loss function that combines multiple criteria: (i) L_1 – hourly MSE: Mean
 squared error between predicted and observed hourly rainfall values; (ii) L_2 – error in reproducing the maximum 1-hour
 precipitation; (iii) L_3 – convectivity penalty, defined only for days with daily total ≥ 20 mm and ratio between the maximum
 300 1-hour precipitation and daily total greater than 0.7, emphasizing accurate reconstruction of convective peaks; (iv) L_4 – error
 in hourly precipitation mean; (v) L_5 – error in hourly precipitation variance; (vi) L_6 – daily cumulative absolute error
 computed between predicted and observed hourly rainfall distributions. L_5 was already introduced by (Oates et al., 2025),
 and L_6 serves the same purpose as the Kullback–Leibler divergence also used by (Oates et al., 2025). Unlike (Bhattacharyya
 et al., 2024), who used a basic MSE loss (L_1), our custom loss explicitly penalizes timing and shape errors in intense rainfall
 305 events, which are critical for flood forecasting and infrastructure design. The model was trained using the Adam optimizer
 (learning rate = 0.001), with early stopping based on validation loss.

4.3 Poisson-Cluster Rainfall Generator

The Poisson-Cluster Disaggregation Method (PCDM) is a stochastic framework designed to disaggregate daily precipitation
 totals into hourly values. The method is based on a simplified Poisson cluster process (Onof et al., 2000), where storm events
 310 (or rain cells) are randomly generated within each day according to a Poisson distribution whose rate parameter is
 conditioned on the daily total precipitation P_d . Each storm cluster has a random start time, duration, and intensity, sampled
 from appropriate distributions (uniform, discrete, and gamma, respectively), and the total hourly series is scaled to match the
 original daily total. The number of intra-day events, N , follows a Poisson distribution with parameter $\lambda = (1.5, P_d / D)$, where
 D is a station-specific scaling divisor. Each event is assigned a start time (random hour of the day), a random duration (1 to 3
 315 hours), and an intensity drawn from a gamma distribution (shape = 2, scale = station-specific parameter). Hourly intensities
 are then aggregated and rescaled to ensure mass conservation, i.e., the sum over 24 hours equals the daily input precipitation
 P_d . PCDM is implemented and applied station-wise using pre-calibrated parameters, D divisor and scale of the gamma
 distribution, obtained via in-site iterative minimization of 1-D Wasserstein distance between observed and simulated
 maximum 1-hour precipitation.

320 The PCDM shares conceptual similarities with widely used Neyman-Scott Rectangular Pulse (NSRP) models. For example,
 both rely on Poisson processes to generate clusters of rainfall events. However, key differences exist. The NSRP model is a
 fully stochastic rainfall generator designed to simulate continuous sequences of rainfall at fine temporal resolution, without
 relying on daily precipitation as a constraint. Rainfall is simulated as a sequence of storms, each consisting of clusters of rain
 325 cells with rectangular pulses. The model requires a larger number of parameters, including those governing storm arrival
 rate, cluster properties, cell duration, and intensity. As such, NSRP is more flexible and capable of reproducing the temporal



structure of rainfall over longer periods, but it is also more complex to calibrate and computationally intensive. In addition, using a gamma distribution in PCDM instead of the classical exponential distribution used in NSRP for rainfall intensity provides greater flexibility and better alignment with observed data. The gamma distribution allows independent control of both the mean and variability, offering improved fit to sub-daily rainfall intensities and better representation of extreme events. This results in more realistic disaggregated series, particularly when preserving hourly maxima is a priority. Another important distinction is that PCDM enforces exact agreement with daily totals, which is advantageous in contexts where observed daily data are available and fidelity to these values is required. In contrast, NSRP achieves mass conservation only in a statistical sense over long simulations, not necessarily for individual days. In summary, while both methods simulate hourly rainfall using event-based structures, PCDM offers a simpler, data-driven alternative to NSRP when the goal is to disaggregate daily precipitation observations to hourly scales, with an emphasis on preserving sub-daily extremes and ensuring mass consistency.

4.4 Multiplicative Microcanonical Random Cascade (MMRC)

Among the family of Multiplicative Microcanonical Random Cascade (MMRC) models, we selected the implementation provided by MELODIST (Förster et al., 2016), as it has been successfully applied in various recent studies (Görner et al., 2021; Hasan et al., 2023; Mayer et al., 2024) and offers a balance between physical realism, probabilistic variability, and computational simplicity. The Cascade Model implemented in MELODIST (MEteoroLOGical observation time series DISaggregation Tool) offers a purely statistical yet physically consistent approach for disaggregating daily precipitation into sub-daily resolutions. Originally proposed by (Olsson, 1998), this MMRC has since been enhanced by subsequent studies to better preserve sub-daily rainfall characteristics while maintaining low computational complexity. The cascade model implemented in MELODIST operates through a hierarchical, stepwise disaggregation in which each level of the time series is recursively split into two subintervals of equal duration. At each level, the precipitation in a given time box is divided into two sub-boxes by assigning a pair of weights (W_1 , W_2), such that $P_{i+1,1} = W_1 \cdot P_i$ and $P_{i+1,2} = (1 - W_1) \cdot P_i = W_2 \cdot P_i$ with the mass-preserving constraint $W_1 + W_2 = 1$. This recursive branching continues until the desired temporal resolution is reached (e.g., 1 hour). For daily to hourly disaggregation, this typically involves the sequence: 24 h $\rightarrow$ 12 h $\rightarrow$ 6 h $\rightarrow$ 3 h $\rightarrow$ 1.5 h $\rightarrow$ 0.75 h $\rightarrow$ 1 h. The model distinguishes between four types of “wet” time boxes: (i) Starting box, follows a dry interval and precedes a wet one; (ii) Ending box, follows a wet interval and precedes a dry one; (iii) Isolated box, surrounded by dry intervals; and (iv) Enclosed box, surrounded by wet intervals. Each of these box types is associated with a set of probabilistic rules for assigning the weights (W_1 , W_2), with three possible branching scenarios: (i) (0, 1), all precipitation goes to the second subinterval; (ii) (1, 0), all precipitation goes to the first subinterval; and (iii) (x , $1 - x$), where precipitation is split proportionally with a random weight $x \in (0, 1)$. The probabilities associated with each case are derived empirically through a reverse scaling procedure applied to high-resolution rainfall data. During this process, fine-resolution data is aggregated up the cascade (e.g., 1 h $\rightarrow$ 2 h $\rightarrow$ 4 h $\rightarrow$ 8 h, etc.), and the branching weights and their frequencies are recorded for each box type and precipitation intensity class (above or below the mean). Instead of fitting parametric distributions, the model uses



empirical histograms of weight values ($x \in [0, 1]$) divided into seven bins. These empirical distributions and associated branching probabilities are stored in lookup tables, which serve as input for a random number generator during disaggregation. For each cascade step during disaggregation, a random number is used to determine the branching type—either (0, 1), (1, 0), or ($x, 1 - x$); if the ($x, 1 - x$) case is selected, a second random number is drawn to choose the specific weight bin x , and the precipitation volume is then split accordingly between the two resulting time steps, ensuring that the total mass is preserved at every level. Because the process is stochastic, different disaggregation realizations may vary, though their statistical properties remain consistent by design. The MMRC model preserves key physical constraints such as precipitation volume conservation (microcanonical property) and captures the intermittency and variability of sub-daily rainfall events with minimal parameterization. The cascade algorithm is a computationally efficient and robust method for producing sub-daily rainfall inputs for hydrological modelling, especially when sub-daily observations are unavailable or sparse.

5 Evaluation Framework

To rigorously evaluate the skill and generalization capacity of the proposed Q-CODA, we implement a 5-fold cross-validation procedure, benchmarking its performance against four established disaggregation methods described in Sect. 4: KNN, ANN-K, PCDM, and MMRC. This comparative framework ensures a systematic and fair evaluation across a spectrum of methodologies. Each fold corresponds to a distinct multi-year period: Fold 1 (1996–2001), Fold 2 (2002–2007), Fold 3 (2008–2013), Fold 4 (2014–2019), and Fold 5 (2020–2024). Each fold is disaggregated using the others for training and finally the 5-folds are joint to reconstruct the whole evaluation period.

The first part of the evaluation focuses on 1-hour maximum precipitation ($P_{1h}^{\max}$), a critical variable for hydrological simulations, risk assessment, and infrastructure design due to its direct relation to peak intensities and convective extremes. By centring the analysis on this high-impact variable, we ensure that the evaluation prioritizes features most relevant to extreme event representation. Model skill is quantified using standard accuracy metrics, Nash–Sutcliffe Efficiency (NSE) and Mean Absolute Error (MAE), calculated across the full $P_{1h}^{\max}$ time series. In addition, distributional similarity is assessed using the first-order Wasserstein distance between observed and simulated $P_{1h}^{\max}$ distributions. To further probe distributional fidelity, we compute relative biases (%) in the mean, variance, and upper-tail quantiles (99th, and 99.9th percentiles), and maximum value of $P_{1h}^{\max}$. Beyond daily extremes, the second component of the evaluation addresses the temporal consistency and structural realism of the full disaggregated hourly time series. Metrics include bias in the temporal mean and temporal variance, the mean duration of precipitation events, as well as the zero proportion (the fraction of dry hours) which captures rainfall intermittency. These structural descriptors are essential for ensuring that the internal dynamics of rainfall sequences are plausible and useful for process-based modelling. Temporal correlation structure is further assessed via biases in the autocorrelation coefficients at lags of 1, 2, 6, and 12 hours, providing insight into the persistence and



organization of precipitation over sub-daily timescales. Finally, to assess the hydrological applicability of each method, we compare Intensity–Duration–Frequency (IDF) curves derived from disaggregated outputs to those obtained from observed hourly records. For the estimation of IDF curves, the classical annual maxima approach was adopted, which remains the standard in hydrological applications. Although alternative methods such as the peaks-over-threshold (POT) framework exist, the annual maxima method was selected to ensure comparability with previous studies, given that the primary focus of this work is the evaluation of the disaggregation algorithm. For each duration (1–24 h), annual maxima were extracted from the hourly precipitation series using a rolling sum. These maxima were fitted to a Gumbel distribution using maximum likelihood estimation. Returns levels (mm) were then computed with the Gumbel inverse CDF:

$$x_T = \mu - \beta \cdot \ln[-\ln(1 - 1/T)], \quad (9)$$

where μ and β are the fitted parameters and T is the return period (years). Intensities (mm h^{-1}) were obtained by dividing each estimate by its duration.

This analysis is conducted at all sites, enabling evaluation of each model’s capacity to reproduce local extremes and temporal scaling behaviours under varied climatic conditions. This multi-faceted validation framework thus offers a robust basis for comparing the accuracy, distributional fidelity, and physical realism of alternative disaggregation approaches.


6 Results

This section presents the results of the 5-fold cross-validation, comparing the performance of Q-CODA against the benchmark disaggregation methods. The aim is to highlight where Q-CODA improves upon existing approaches and to identify trade-offs across the evaluated metrics.

Figure 5 presents a comparative evaluation of the disaggregation methods based on the ability to reconstruct the daily maximum 1-hour precipitation ($P_{1h}^{\max}$) across 91 meteorological stations. Each panel displays boxplots that summarize the spatial distribution of different performance metrics for each method. The evaluated metrics include: Mean Absolute Error (MAE) (Fig. 5a), Nash–Sutcliffe Efficiency (NSE) (Fig. 5b), 1-D Wasserstein distance (Fig. 5c), bias in the mean (Fig. 5d), variance (Fig. 5e), 99th percentile (Fig. 5f), 99.9th percentile (Fig. 5g), and daily maximum (Fig. 5h). Across all these metrics, Q-CODA consistently outperforms all benchmark methods, demonstrating superior accuracy and robustness in reproducing high-resolution rainfall extremes. In terms of MAE (Fig. 5a), Q-CODA exhibits the lowest errors across the station network, followed closely by ANN-K. The remaining methods—KNN, PCDM, and MMRC—yield significantly higher errors, with highly similar performance distributions. For the NSE (Fig. 5b), Q-CODA and ANN-K attain the highest median values, approaching 0.75, while KNN, PCDM, and MMRC display lower medians in the range of 0.5 to 0.6. Notably, MMRC exhibits a broader interquartile range and a substantially extended lower whisker, indicating greater spatial variability and underperformance at some locations. The results for the 1-D Wasserstein distance (Fig. 5c) further confirm



the advantage of Q-CODA, which yields the lowest values overall, indicating the closest alignment with the observed distribution of $P_{1h}^{\max}$. KNN ranks second, while the other methods perform markedly worse. Interestingly, ANN-K performs the worst despite its custom loss function, underscoring the limitations of neural networks in preserving distributional shapes under regression pressure. With respect to the bias in the mean (Fig. 5d), Q-CODA again leads with minimal bias, followed by KNN. In contrast, ANN-K and MMRC exhibit wide interquartile ranges, and PCDM shows a consistent negative bias. The evaluation of variance bias (Fig. 5e) highlights a major strength of Q-CODA, which substantially outperforms the alternatives. KNN and MMRC tend to overestimate variance, whereas ANN-K and PCDM show systematic underestimation, a well-documented limitation of neural network-based disaggregation (Oates et al., 2025; Bhattacharyya et al., 2024). Regarding extreme rainfall reconstruction, panels 5f through 5h show biases in the 99th percentile, 99.9th percentile, and maximum of daily $P_{1h}^{\max}$ values. Q-CODA delivers remarkably low biases in the 99th and 99.9th percentiles and remains the best-performing method even when evaluating the absolute maximum, despite a slight increase in error. The other methods struggle considerably with tail behaviour: ANN-K systematically underestimates extremes, MMRC tends to overestimate them, and PCDM performs similarly to ANN-K in this regard. KNN offers moderately accurate results but falls short of Q-CODA. The poor performance of ANN-K in reproducing maximum $P_{1h}^{\max}$ may also stem from the well-known limitation of deep learning models to extrapolate beyond the range of their training data.



440 **Figure 5: Comparative evaluation of disaggregation methods based on their ability to reconstruct daily maximum 1-hour precipitation ($P_{1h}^{\max}$) across 91 meteorological stations. Each panel shows boxplots summarizing the spatial distribution of key performance metrics: (a) Mean Absolute Error (MAE), (b) Nash–Sutcliffe Efficiency (NSE), (c) 1-D Wasserstein distance, (d) bias in the mean, (e) bias in the variance, (f) bias in the 99th percentile, (g) bias in the 99.9th percentile, and (h) bias in the maximum.**



Overall, the results in Fig. 5 demonstrate the superior ability of Q-CODA to reproduce both the magnitude and statistical characteristics of peak sub-daily precipitation. Its consistent performance across multiple metrics and spatial locations highlights its value as a robust and accurate disaggregation tool.

Figure 6 presents a comparative assessment of the disaggregation methods based on their ability to reproduce key statistical and temporal properties of the full hourly time series at the 91 evaluation stations. Each subplot shows the spatial distribution of bias (%) for different metrics, summarized through boxplots, where each box reflects variability across sites. The analysed metrics include: bias in the temporal mean (Fig. 6a), temporal variance (Fig. 6b), zero proportion (Fig. 6c), mean event duration (Fig. 6d), and autocorrelation at lag-1 (Fig. 6e), lag-2 (Fig. 6f), lag-6 (Fig. 6g), and lag-12 (Fig. 6h). For the mean of the hourly series (Fig. 6a), Q-CODA exhibits an almost negligible bias, clearly outperforming the other methods. KNN, PCDM, and MMRC follow with slight underestimations, while ANN-K stands out negatively, displaying a more pronounced underestimation close to 0.5%. The results for the variance (Fig. 6b) show a similar pattern. Q-CODA again achieves the lowest bias, followed by KNN, PCDM, and MMRC, which tend to modestly underestimate variability. ANN-K shows the most significant underestimation, with a median bias exceeding 25%, confirming the well-known difficulty of neural networks in reproducing precipitation variability, as previously documented by (Bhattacharyya et al., 2024; Oates et al., 2025). When evaluating the proportion of zero values (Fig. 6c), which corresponds to the share of dry hours, KNN achieves the lowest bias, followed closely by Q-CODA and MMRC. ANN-K and PCDM show worse performance, with ANN-K underestimating and PCDM overestimating the frequency of dry periods. For the mean event duration (Fig. 6d), Q-CODA again demonstrates the best performance, yielding the smallest bias across stations. KNN and PCDM both tend to underestimate event lengths, while MMRC shows a tendency to overestimate. ANN-K has its median bias near zero, but a large interquartile range indicates high spatial variability and reduced robustness. Autocorrelation results further highlight Q-CODA's superiority. In the lag-1 (Fig. 6e), lag-2 (Fig. 6f), lag-6 (Fig. 6g), and lag-12 (Fig. 6h) cases, Q-CODA consistently provides the lowest median bias. While MMRC shows the smallest median bias at lag-12, Q-CODA achieves a narrower interquartile range, suggesting greater spatial consistency. KNN systematically underestimates autocorrelation across all lags. PCDM and MMRC also show underestimation, except for lag-12 where MMRC slightly improves. The issue of autocorrelation underestimation in the MMRC approach had previously been identified and reported by (Müller and Haberlandt, 2018). ANN-K transitions from underestimation at lag-1, to a nearly unbiased result at lag-2, but then shifts to strong overestimation at lag-6 and lag-12, pointing to instability in its representation of temporal structure.

Overall, the metrics in Fig. 6 underscore the robustness and generalization ability of Q-CODA. It achieves the best or near-best performance across all statistical and temporal features of the hourly rainfall series, offering a consistent advantage over benchmark methods.




Figure 6. Comparative evaluation of disaggregation methods based on their ability to reconstruct complete hourly precipitation time series across 91 meteorological stations. The panel presents boxplots summarizing the spatial distribution of key performance metrics: (a) mean hourly precipitation, (b) hourly precipitation variance, (c) bias in zero proportion, (d) bias in mean precipitation event duration, (e) bias in lag-1 autocorrelation, (f) bias in lag-2 autocorrelation, (g) bias in lag-6 autocorrelation, and (h) bias in lag-12 autocorrelation.



The realistic reproduction of key temporal rainfall characteristics, such as the proportion of zero values, the mean precipitation event duration, and the autocorrelation structure, is critical in hydrological applications, as these attributes directly influence runoff generation, soil moisture dynamics, and flood modelling. An over- or underestimation of dry intervals (i.e., zero proportion) can distort the temporal distribution of rainfall, while inaccuracies in event duration may lead to erroneous estimates of infiltration and surface flow. Similarly, the autocorrelation of rainfall, particularly at short lags, affects the temporal persistence of wet conditions, which in turn shapes hydrological responses at various scales. Among the evaluated disaggregation methods, Q-CODA stands out by most accurately preserving these features across the network of stations. Its ability to replicate the observed structure of rainfall time series with high fidelity enhances the reliability of downstream hydrological simulations, making it a robust tool for both historical analysis and future climate impact assessments.

To evaluate the ability of each disaggregation method to reproduce observed IDF (Intensity–Duration–Frequency) relationships, we conducted an illustrative case study at the Reus station (41.1450°N, 1.1636°E), which lies within a Mediterranean climate zone (Köppen-Geiger climate classification Csa). This site was selected because it exhibits the highest recorded 1-hour precipitation among all stations in the historical dataset. Figure 7 compares the IDF curves obtained from each method against the observed benchmark for return periods of 2, 5, 10, 25, 50, and 100 years. Q-CODA closely tracks the empirical IDF across all return periods, exhibiting minimal deviation even at more extreme thresholds. MMRC also performs well in this case, yielding simulated IDF curves that remain relatively consistent with the observations. By contrast, the IDF curves produced by KNN, ANN-K, and PCDM deviate more substantially from the reference, especially at longer return periods. While informative, this example represents a single location and does not capture the broader spatial variability in IDF performance. To address this, we computed the root mean squared error (RMSE) between simulated and observed IDF curves at each station and for each method, yielding a spatial distribution of RMSE values summarized via boxplots in Fig. 8. This analysis reveals that Q-CODA consistently achieves the lowest IDF reconstruction error across the network of 91 stations, followed (at increasing distances) by KNN, MMRC, PCDM, and ANN-K. This ranking remains stable across all return periods, although RMSE values increase with return period, reflecting the greater challenge of accurately modelling rare extremes. These results underscore Q-CODA’s superior capacity to preserve high-order rainfall statistics essential for hydrological design and risk assessment.




Figure 7. Comparison of Intensity–Duration–Frequency (IDF) curves derived from each disaggregation method with the observed benchmark at the Reus station (41.1450°N , 1.1636°E). Curves are shown for return periods of 2, 5, 10, 25, 50, and 100 years.




Figure 8. Spatial distribution of root mean squared error (RMSE) between observed and simulated IDF (Intensity–Duration–Frequency) curves across 91 meteorological stations, grouped by disaggregation method. Boxplots summarize the performance variability of each method in reproducing IDF relationships, providing a broader evaluation beyond the single-station case study shown in Fig. 7.

520



7 Conclusions

This study introduced Q-CODA, a novel data-driven rainfall disaggregation framework that leverages the quasi-comonotonic relationship between daily precipitation totals and sub-daily maxima. Building on this empirically supported dependence structure, Q-CODA, through a comonotonic transformation, ensures that the simulated sub-daily maxima lie within plausible bounds conditional on the daily total, especially preserving the upper tail behaviour essential for extreme event modelling. Subsequently, Q-CODA initiates the disaggregation process by seeding an initial hourly pattern via a seasonal K-Nearest Neighbours (KNN) search, selecting and scaling realistic historical analogs. This pattern is then iteratively refined to satisfy both the daily total and the set of comonotonicity-based sub-daily maxima constraints across multiple durations (1 h, 2 h, 6 h, and 12 h). The adjustment procedure further enhances the temporal structure by increasing lag-1 autocorrelation through targeted reordering and smoothing, ensuring realism in the resulting rainfall sequences. This hybrid approach combines probabilistic rigor, empirical realism, and interpretability.

A comprehensive benchmarking against four existing methods—KNN, ANN-KMeans (ANN-K), Microcanonical Multiplicative Random Cascade (MMRC), and the Poisson Cluster Disaggregation Model (PCDM)—demonstrates that Q-CODA consistently outperforms alternatives across a wide range of metrics. These include deterministic accuracy (e.g., MAE, NSE), distributional fidelity (bias in mean, variance, and high quantiles), and temporal structure (autocorrelation, event duration, dry hour frequency), evaluated across 91 stations in Spain. In particular, Q-CODA excels at reproducing extremes such as the 99th and 99.9th percentiles of hourly maxima and produces the lowest error in simulated IDF curves.

The comparative performance of ANN-K highlights some benefits of deep learning in rainfall disaggregation, but also reveals important drawbacks, such as underestimation of variance and systematic misrepresentation of extremes, partially due to the limitations of neural networks in extrapolation beyond the training range. Traditional approaches such as MMRC and PCDM capture some statistical properties but show greater variability and biases across stations. Overall, Q-CODA emerges as a robust, flexible, and physically grounded methodology for disaggregating daily rainfall into sub-daily time series. Its explicit use of dependence structures via comonotonic transformations, combined with an efficient constraint-based adjustment scheme, makes it well suited for hydrological modelling, flood risk assessment, and climate adaptation planning because it can be used for the extension of insufficient hourly precipitation data or the temporal disaggregation of climate change daily projections.

Future research may explore several promising directions. Extending Q-CODA to include spatial coherence across station networks or gridded products would enable its application in basin-scale hydrological simulations. Coupling the methodology with regional climate model outputs could allow integration into broader downscaling frameworks. Moreover, increasing the temporal resolution from hourly to sub-hourly or minute-scale disaggregation is particularly relevant for high-



555 resolution hydrological applications, including urban drainage, flash flood forecasting, and early warning systems, where rainfall dynamics at minute-level resolution are critical.

Code availability

560 The source code used to implement the Q-CODA disaggregation method is openly available at GitHub (<https://github.com/carloscorreag/pyqcoda>) and has been archived in Zenodo with the DOI <https://doi.org/10.5281/zenodo.16364100>. Additionally, the package is distributed via the Python Package Index (PyPI) at <https://pypi.org/project/pyqcoda/>, enabling straightforward installation through standard Python package managers.

Author contribution

565 C. Correa: Conceptualization; Data curation; Formal analysis; Methodology; Software; Validation; Visualization; and Writing (original draft preparation).

A. Hernanz: Conceptualization; Data curation; Writing (review and editing).

I. San-Felipe: Data curation.

E. Rodríguez-Guisado: Writing (review and editing).

Competing interests

570 The authors declare that they have no conflict of interest.

Acknowledgements

The authors gratefully acknowledge the National Climate Data Bank of AEMET (Agencia Estatal de Meteorología, Spain) for providing access to the high-quality hourly precipitation records used in this study.

References

- 575 Abdellatif, M., Kuchling, P., Rüdiger, B., and Ventura, I.: Wasserstein distance in terms of the comonotonicity copula, *Stochastics*, 1–12, doi:10.1080/17442508.2024.2427734, 2024.
- Alzahrani, F., Seidou, O., and Alodah, A.: Assessing the performance of daily to subdaily temporal disaggregation methods for the IDF curve generation under climate change, *J. Water Clim. Change*, 14, 1339–1357, doi:10.2166/wcc.2023.507, 2023.



- 580 Bhattacharyya, D., and Saha, U.: Deep learning application for disaggregation of rainfall with emphasis on preservation of extreme rainfall characteristics for Indian monsoon conditions, *Stoch. Environ. Res. Risk Assess.*, 37, 1021–1038, doi:10.1007/s00477-022-02331-x, 2023.
- Bhattacharyya, D., Deka, P., and Saha, U.: Applicability of statistical and deep-learning models for rainfall disaggregation at metropolitan stations in India, *J. Hydrol. Reg. Stud.*, 51, 101616, doi:10.1016/j.ejrh.2023.101616, 2024.
- 585 Biswas, P., and Saha, U.: Disaggregation of rainfall from daily to 1-hour scale through integrated MMRC-copula modelling, *J. Hydrol.*, 647, 132338, doi:10.1016/j.jhydrol.2024.132338, 2025.
- Chakravarti, N.: Isotonic median regression: a linear programming approach, *Math. Oper. Res.*, 14, 303–308, doi:10.1287/moor.14.2.303, 1989.
- De Leeuw, J.: Correctness of Kruskal's algorithms for monotone regression with ties, *Psychometrika*, 42, 141–144, doi:10.1007/BF02293750, 1977.
- 590 Delong, L. and Wüthrich, M. V.: Isotonic regression for variance estimation and its role in mean estimation and model validation, *North Am. Actuar. J.*, 1–29, doi:10.1080/10920277.2024.2421221, 2024.
- Ebers, N., Schröter, K., and Müller-Thomy, H.: Estimation of future rainfall extreme values by temperature-dependent disaggregation of climate model data, *Nat. Hazards Earth Syst. Sci.*, 24, 2025–2043, doi:10.5194/nhess-24-2025-2024, 2024.
- 595 Förster, K., Hanzer, F., Winter, B., Marke, T., and Strasser, U.: An Open-Source MEteoroLOGical Observation Time Series DISaggregation Tool (MELODIST v0.1.1), *Geosci. Model Dev.*, 9, 2315–2333, doi:10.5194/gmd-9-2315-2016, 2016.
- Fowler, H. J., Wasko, C., and Prein, A. F.: Intensification of short-duration rainfall extremes and implications for flood risk: Current state of the art and future directions, *Philos. Trans. R. Soc. A*, 379, 20190541, doi:10.1098/rsta.2019.0541, 2021.
- Görner, C., Franke, J., Kronenberg, R., Hellmuth, O., and Bernhofer, C.: Multivariate non-parametric Euclidean distance model for hourly disaggregation of daily climate data, *Theor. Appl. Climatol.*, 143, 241–265, doi:10.1007/s00704-020-03426-7, 2021.
- 600 Hasan, M. M., Strong, C., Brooks, P. D., Burian, S. J., and Barber, M. E.: Quantifying climate change impacts on low flows of small high mountain watersheds: A nonstationary approach, *J. Hydrol. Reg. Stud.*, 48, 101463, doi:10.1016/j.ejrh.2023.101463, 2023.
- 605 Haslinger, K., Breinl, K., Pavlin, L., Pistotnik, G., Bertola, M., Olefs, M., Greilinger, M., Schöner, W., and Blöschl, G.: Increasing hourly heavy rainfall in Austria reflected in flood changes, *Nature*, 639, 667–672, doi:10.1038/s41586-025-08647-2, 2025.
- Joe, H.: Asymptotic efficiency of the two-stage estimation method for copula-based models, *J. Multivar. Anal.*, 94, 401–419, doi:10.1016/j.jmva.2004.06.003, 2005.
- 610 Li, X., Meshgi, A., Wang, X., Zhang, J., Tay, S. H. X., Pijcke, G., and Babovic, V.: Three resampling approaches based on method of fragments for daily-to-subdaily precipitation disaggregation, *Int. J. Climatol.*, 38, e1119–e1138, doi:10.1002/joc.5438, 2018.



- Lopez-Gomez, I., McGovern, A., Agrawal, S., and Hickey, J.: Global extreme heat forecasting using neural weather models, *Artif. Intell. Earth Syst.*, 2(1), e220035, doi:10.1175/AIES-D-22-0035.1, 2023.
- 615 Maloku, K., Hingray, B., and Evin, G.: Accounting for precipitation asymmetry in a multiplicative random cascade disaggregation model, *Hydrol. Earth Syst. Sci.*, 27, 3643–3661, doi:10.5194/hess-27-3643-2023, 2023.
- Mayer, S., Hendrick, M., Michel, A., Richter, B., Schweizer, J., Wernli, H., and van Herwijnen, A.: Impact of climate change on snow avalanche activity in the Swiss Alps, *The Cryosphere*, 18, 5495–5517, doi:10.5194/tc-18-5495-2024, 2024.
- Moradzadeh Rahmatabadi, S., Irandoust, M., and Mirabbasi, R.: Multivariate analysis of rainfall–runoff characteristics using
 620 copulas, *J. Earth Syst. Sci.*, 132, 93, doi:10.1007/s12040-023-02105-1, 2023.
- Müller, H., and Haberlandt, U.: Temporal rainfall disaggregation using a multiplicative cascade model for spatial application in urban hydrology, *J. Hydrol.*, 556, 847–864, doi:10.1016/j.jhydrol.2016.01.031, 2018.
- Müller-Thomy, H.: Temporal rainfall disaggregation using a micro-canonical cascade model: possibilities to improve the autocorrelation, *Hydrol. Earth Syst. Sci.*, 24, 169–188, doi:10.5194/hess-24-169-2020, 2020.
- 625 Naderi, K., Moghaddasi, M., and Shokri, A.: Drought occurrence probability analysis using multivariate standardized drought index and copula function under climate change, *Water Resour. Manag.*, 36, 2865–2888, doi:10.1007/s11269-022-03186-1, 2022.
- Oates, H., Arora, N., Oh, H. G., and Lee, T.: A long short-term memory model for sub-hourly temporal disaggregation of precipitation, *Stoch. Environ. Res. Risk Assess.*, 1–14, doi:10.1007/s00477-025-02996-0, 2025.
- 630 Olsson, J.: Evaluation of a scaling cascade model for temporal rainfall disaggregation, *Hydrol. Earth Syst. Sci.*, 2, 19–30, doi:10.5194/hess-2-19-1998, 1998.
- Onof, C., Chandler, R. E., Kakou, A., Northrop, P., Wheeler, H. S., and Isham, V.: Rainfall modelling using Poisson-cluster processes: a review of developments, *Stoch. Environ. Res. Risk Assess.*, 14, 384–411, doi:10.1007/s004770000043, 2000.
- Pui, A., Sharma, A., Mehrotra, R., Sivakumar, B., and Jeremiah, E.: A comparison of alternatives for daily to sub-daily
 635 rainfall disaggregation, *J. Hydrol.*, 470, 138–157, doi:10.1016/j.jhydrol.2012.08.041, 2012.
- Qin, X., and Dai, C.: Investigation of rainfall disaggregation with flexible timescales based on point process models, *J. Hydrol.*, 634, 131101, doi:10.1016/j.jhydrol.2024.131101, 2024.
- Rodriguez-Iturbe, I., Cox, D. R., and Isham, V.: Some models for rainfall based on stochastic point processes, *Proc. R. Soc. Lond. A Math. Phys. Sci.*, 410, 269–288, doi:10.1098/rspa.1987.0039, 1987.
- 640 Rousseeuw, P. J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis, *J. Comput. Appl. Math.*, 20, 53–65, doi:10.1016/0377-0427(87)90125-7, 1987.
- Scher, S., and Peßenteiner, S.: Temporal disaggregation of spatial rainfall fields with generative adversarial networks, *Hydrol. Earth Syst. Sci.*, 25, 3207–3225, doi:10.5194/hess-25-3207-2021, 2021.
- Schilling, W.: Rainfall data for urban hydrology: what do we need?, *Atmos. Res.*, 27, 5–21, doi:10.1016/0169-
 645 8095(91)90003-F, 1991.



- Sklar, A.: Random variables, distribution functions, and copulas—a personal look backward and forward, in Distributions with Fixed Marginals and Related Topics (Seattle, WA, 1993), IMS Lecture Notes Monogr. Ser., 28, Inst. Math. Statist., Hayward, CA, 1–14, doi:10.1214/lnms/1215452606, 1996.
- Stoner, O., and Economou, T.: An advanced hidden Markov model for hourly rainfall time series, Comput. Stat. Data Anal., 152, 107045, doi:10.1016/j.csda.2020.107045, 2020.
- Takhellambam, B. S., Srivastava, P., Lamba, J., McGehee, R. P., Kumar, H., and Tian, S.: Temporal disaggregation of hourly precipitation under changing climate over the Southeast United States, Sci. Data, 9, 211, doi:10.1038/s41597-022-01304-7, 2022.
- Vorobevskii, I., Park, J., Kim, D., Barfus, K., and Kronenberg, R.: Simulating sub-hourly rainfall data for current and future periods using two statistical disaggregation models: case studies from Germany and South Korea, Hydrol. Earth Syst. Sci., 28, 391–416, doi:10.5194/hess-28-391-2024, 2024.