

Probabilistic Rainfall Downscaling: Joint Generalized Neural Models with Censored Spatial Gaussian Copula

David Huk

*Department of Statistics
University of Warwick
Coventry, CV4 7AL
United Kingdom*

DAVID.HUK@WARWICK.AC.UK

Rilwan A. Adewoyin *

*Department of Statistics
University of Warwick
Coventry, CV4 7AL
United Kingdom*

RILWAN.ADEWOYIN@WARWICK.AC.UK

Ritabrata Dutta

*Department of Statistics
University of Warwick
Coventry, CV4 7AL
United Kingdom*

RITABRATA.DUTTA@WARWICK.AC.UK

Editor: Arindam Banerjee

Abstract

A novel approach for generating conditional probabilistic rainfall downscaling at finer scales from deterministic weather variables at coarser scales with temporal and spatial dependence is introduced. A two-step procedure is employed. Firstly, marginal location-specific distributions are jointly modelled conditional on the deterministic coarse weather variables. Secondly, a spatial dependency structure is learned to ensure spatial coherence among these distributions. To learn marginal distributions over rainfall values, we introduce joint generalised neural models that expand generalised linear models with a deep neural network architecture to jointly fit parameters of the distributions. The spatial dependency structure is modelled using a censored latent Gaussian copula leveraging the underlying spatial structure. We construct a distance matrix between locations, transformed into a correlation matrix by a Gaussian Process Kernel depending on a small set of parameters. To estimate these parameters, we propose a general framework for the estimation of latent Gaussian copulas employing scoring rules as a measure of divergence between distributions. Uniting our two contributions, namely the joint generalised neural model and the censored latent Gaussian copulas into a single model, our probabilistic approach provides downscaled rainfall. We demonstrate its efficacy using a large UK data set, outperforming existing methods.

Keywords: Joint Generalised Neural Model; Censored Gaussian copula; Precipitation downscaling; Spatio-temporal forecasting; Minimum scoring rules

*. Also affiliated with Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China.

1. Introduction

In an era of escalating climate change, the global community is witnessing an intensifying frequency of high precipitation events. As per the United Nations’ report, flooding emerged as the most prevalent weather-induced disaster over the two decades culminating in 2015, impacting 2.3 billion individuals and resulting in an economic toll of approximately 1.89 trillion (Wallemacq and Herden, 2015). In response to the mounting fiscal and societal risks associated with heavy precipitation and flooding, policy makers are increasingly leveraging large-scale numerical weather predictors (NWP) (Randall et al., 2007; Jung et al., 2010) in conjunction with flood models (FMs) (Qi et al., 2021).

One of the main shortcomings of NWP is their coarse resolution. State-of-the-art NWP act on horizontal scales of 50 km to 100 km, which is insufficient to capture local precipitation. But to accurately predict possible future flooding scenarios, FMs do need an accurate prediction of precipitation on a fine grid (for example, a 10 km grid) for all locations simultaneously in a large area (such as all of the UK). To provide local-scale precipitation predictions, it is common practice to downscale coarse NWP forecasts to a finer grid with the use of post-processing methods. Similar to the works by Woo and Wong (2017); Vandal et al. (2018); Adewoyin et al. (2021); Xiang et al. (2022) and Wang et al. (2023), one of our main aims is to perform statistical *downscaling* of rainfall, that is predicting high-resolution precipitation from low-resolution weather variables.

Furthermore, one would like to have uncertainty quantification about possible rainfall levels over a large area in order to assess various precipitation-related risks (e.g. risks of different flooding scenarios) (Schoof, 2013), which is not possible with deterministic downscaling, motivating our model to provide *probabilistic* downscaling of precipitation. Predicting a distribution for rainfall rather than merely predicting the expected value can also rectify the deficiency in the simulation of extreme precipitation (Watson, 2022), as point estimates are inherently biased towards the more commonly observed low rainfall events. To provide accurate estimates of risks of different flooding scenarios through flood models, we need to be able to predict and sample from the joint distribution of precipitation from a large area in a *spatially-coherent* manner rather than only providing probabilistic prediction of precipitation in a location-specific manner. Unlike implicit generative Machine Learning based approaches such as Ravuri et al. (2021); Harris et al. (2022); Chen et al. (2024); Duncan et al. (2022) and Price and Rasp (2022) which do not have access to the distribution or cannot output precisely zero-valued observations, we would like to obtain an explicit model for rainfall quantities capable of predicting exact zero amounts, thus retaining the interpretability of the model’s outputs.

Hence our **main aim** is the development of a downscaling method for the reliable and explainable mapping of large-scale atmospheric variables from coarse weather or climate models to *spatially-coherent high-resolution probabilistic predictions of precipitation*.

To this aim, we introduce a pioneering two-component probabilistic model for the downscaling of rainfall. The first component comprises a novel joint generalised neural model designed to learn location-specific explicit marginal distributions conditioned on spatio-temporal covariates through a novel architecture. (Note that here, “joint” refers to the estimation of marginal density parameters being done jointly, not to be confused with the estimation of a joint density). The second component is a censored latent Gaussian copula

approach, specifically designed for censored data, which introduces spatial dependency into the location-specific downscaling of precipitation. When combined, these two components generate spatially and temporally coherent joint distributions of rainfall. We provide a diagram of our model, termed Cens-JGNM, in Figure 1. We apply our model to an extensive UK rainfall data set and juxtapose its performance with established benchmark methods. Our approach has important applications, including:

- The extension of univariate forecasts to joint distributions with coherent spatial dependence. Many models in the Geoscience community exist for reliable location-wise prediction of censored observations. Our methodology offers an easy extension of the existing models to simultaneously predict variables of interest at multiple stations while taking into account both the censoring of data and the inherent spatial dependency.
- The improvement of rainfall predictions (for single or multiple locations) at a local level by enhancing their resolution with our approach, taking additional high-resolution information into account. Rainfall predictions are notably more challenging and have a shorter prediction horizon compared to other atmospheric variables due to the highly variable and localized nature of precipitation events. Factors such as complex cloud microphysics, rapid changes in atmospheric moisture, and small-scale convective processes contribute to this decreased predictability as they evolve at a scale significantly finer than our input data. Consequently, our approach aims to extend the temporal reach of rainfall forecasts beyond that of standard weather or climate models by exploiting correlations between more predictable atmospheric features and precipitation. Furthermore, our explicit modelling of distributions ensures explainability and captures the uncertainty inherent in precipitation downscaling.
- The computationally efficient generation of downscaled total precipitation outputs that enhance the accuracy of total precipitation predictions from leading NWP models, such as ERA5, which are known to exhibit significant bias (Lavers et al., 2022). This problem also extends to data driven approaches which rely on NWP output as training data and are prone to learning the biases present. As such, prominent data driven approaches have excluded total precipitation from outputs or diagnostics in (Lang et al., 2024; Lam et al., 2023).

The remainder of this paper is structured as follows: Section 2 introduces the UK meteorological and rainfall data set employed in this study (ERA5 and E-OBS for training and IFS for benchmarking) along with our two-component approach. Section 3 details the application of joint generalised neural models as a method for learning marginal distributions. Section 4 explains how censored latent Gaussian copulas can extend marginal densities into a joint prediction with the correct spatial dependence in the presence of censored observations. Section 5 showcases the results of censored Gaussian copulas fitted with the joint generalised neural model marginal distributions to the UK data set. Connections with existing work are established in Section 6 with this study’s conclusion presented in Section 7.

2. Data and Statistical Problem

We begin by providing a description of the data in Section 2.1, followed by an explanation of our statistical problem and approach in Section 2.2. Notation for the main text is included in Table 1 with deviations from this indicated in the relevant Sections.

Symbol	Description	Symbol	Description
\mathcal{I}	Set of location $\{1, \dots, n\}$	i	Individual location
\mathcal{T}	Set of days	t	Individual day
\mathcal{Y}_s	Target rainfall over all locations in \mathcal{I} for day s	$y_{i,s}$	Target rainfall at location i and day s
\mathcal{Y}_s^*	Latent variable vector over all locations \mathcal{I} for day s	$y_{i,s}^*$	Latent variable at location i and day s
$\mathcal{X}_{i,s}$	Weather variables relevant for location i and day s	$\mathbf{x}_{i,s}$	Subset of predictors $\subseteq \mathcal{H}_s$ relevant to predict $y_{i,s}$
\mathcal{H}_s	History of $\mathcal{X}_{i,s}$ over all i , and for days relevant to day s	\mathcal{D}	Distance matrix $\in \mathbb{R}^{n \times n}$ of spatial information for all locations.
$F_{i,s}$	(Conditional) distribution of rain at location i for day s	$F_{i,s}^*$	Latent (conditional) distribution at location i for day s
$f_{i,s}$	(Conditional) density of rain at location i for day s	$f_{i,s}^*$	Latent (conditional) density at location i for day s
π	Joint (conditional) rainfall density over \mathcal{I} for day s	π^*	Latent joint (conditional) density over \mathcal{I} for day s .
\mathbf{c}	Gaussian copula density in n dimensions	\mathbf{c}^*	Latent copula density of quantiles from $F_{i,s}^*$ over \mathcal{I}
Φ	Standard Gaussian distribution function	ϕ	Standard Gaussian density function
$z_{i,s}$	Observed data on Gaussian scale $\Phi^{-1}(F_{i,s}(y_{i,s} \mathcal{H}_s))$	$z_{i,s}^*$	Latent data on Gaussian scale $\Phi^{-1}\left(F_{i,s}^*(y_{i,s}^* \mathcal{H}_s)\right)$

Table 1: Notations used throughout the main text.

2.1 Meteorological Data

ERA5 and E-OBS data: Our model aims to downscale weather variables from the ERA5 reanalysis data set (Hersbach et al., 2020) on a coarse grid to a joint distribution of rainfall across multiple locations on a finer grid as given by the E-OBS rainfall data (Cornes et al., 2019). Table 2 summarizes the data sources involved in this work.

The ECMWF employs a physics based weather simulator named the Integrated Forecasting System (IFS) for operational weather predictions and to generate the ERA5 (Hersbach et al., 2020) reanalysis data, which serves as input data for our work. We use a selection

of model fields, which can be resolved reasonably well by typical climate models, from the ERA5 reanalysis dataset (Hersbach et al., 2020) as input variables for our model, as we hope to extend our model to climate model inputs in future works. Formally, at 6-hourly intervals $j \in \{1, 2, 3, 4\}$ on day $t \in \mathcal{T} \subset \mathbb{N}$ and for locations $i \in \{1, \dots, n\} = \mathcal{I}$, we are given predictors

$$\mathcal{X}_{i,jt} = (\mathcal{X}_{i,jt}^{(1)}, \dots, \mathcal{X}_{i,jt}^{(6)}),$$

representing six weather variables from the ERA5 reanalysis: specific humidity, air temperature, geopotential height (500 hPa), the eastward and northward components of wind velocity (850 hPa), and total column water vapour. These variables are available four times per day at a resolution of approximately $5/9^\circ$ in longitude and $5/6^\circ$ in latitude (corresponding to roughly 50 km grid spacing over Wales) as produced by the IFS-based reanalysis.

For each day, daily means of these smoothed model fields are computed. However, because our outcome variables—the 24-hour precipitation accumulations (in millimetres) from the E-OBS dataset—are recorded on a 0.1° longitude-latitude grid, it is necessary to interpolate the input fields to the same high-resolution grid. In practice, the original coarse grid (which covers a spatial extent roughly corresponding to a 20×21 grid) is regridded using a bivariate spline interpolation scheme on a rectangular mesh to produce a finer 100×140 grid with 0.1° resolution (approximately 8.5 km spatial resolution). This constitutes our input weather variables $\mathcal{X}_{i,s}, i \in \mathcal{I}, s \in \mathcal{T}$ used for modelling. This interpolation not only refines the resolution but, due to a cropping of the overall extent, ensures that the predictor and outcome variables are spatially aligned.

IFS Benchmark data: As a deterministic benchmark, we use precipitation predictions available from short-term forecast simulations with IFS as they were performed to provide the background state of the 4DVar data-assimilation when ERA5 was generated. There are two forecast simulations started each day at 6 am and 6 pm. We extract the precipitation fields for the first 12 hours of each simulation to reproduce daily precipitation - this is presently the optimal way to derive meaningful precipitation predictions from a dynamical model that is consistent with the large-scale fields in the ERA5 reanalysis data. Note that the IFS precipitation predictions provide only a deterministic prediction for each point in space and time, which requires simulating the full three-dimensional state of the global atmosphere, and are computationally expensive. In contrast, our mapping provides a probabilistically downscaled rainfall value, requiring only 6 local model fields as input, and can be generated at a much lower cost once the model is trained. Furthermore, the precipitation data was mapped directly from the native Gaussian grid of ERA5 (N320) to the high-resolution grid (0.1° longitude and latitude) of the observations and was not mapped to the coarse resolution grid at $5/9^\circ$ longitude and $5/6^\circ$ latitude first. The IFS precipitation predictions have therefore a higher spatial resolution (~ 30 km) when compared to the model fields that serve as input to our model, making it an especially challenging benchmark.

2.2 Statistical Problem

Our goal is to downscale rainfall over the UK by conditioning on weather predictors while portraying uncertainty and preserving spatial coherence. We are presented with predictors in the form of weather variables $\mathcal{X}_{i,t}$ for days $t \in \mathcal{T}$ and location indices i from the set \mathcal{I} of

Source	Weather Variables	Description
ERA5 Dataset	Air Temperature, Specific Humidity, Geopotential Height (500 hPa), Wind Velocity (U/V, 850 hPa), Total Column Water Vapour	Historical weather state estimates from a numerical weather model with continuous observation assimilation. Provides 6-hourly records at 1-degree resolution. Used as inputs for our neural network models.
E-OBS Dataset	Rainfall Daily Total Precipitation	Observational records collected from a network of weather stations. Provides daily records at 0.1-degree resolution. Used as the target dataset for our neural network models.
IFS Precipitation Dataset for benchmarking	Daily Total Precipitation	Forecast precipitation estimates from the free-running IFS model. Unlike ERA5, these forecasts rely solely on internal model physics without continuous assimilation of precipitation observations. Used as a deterministic baseline for comparison.

Table 2: Meteorological Datasets.

locations considered. As we want to provide a probabilistic prediction at a time $s \in \mathcal{T}$ for rainfall $\mathcal{Y}_s = (y_{1,s}, \dots, y_{n,s})$ simultaneously over all locations in \mathcal{I} , we need to construct a conditional joint distribution of rainfall over the n locations at time s , written $\mathbb{P}(\mathcal{Y}_s | \mathcal{H}_s)$, conditioning on the history of predictor variables until time s denoted as $\mathcal{H}_s \equiv \{\mathcal{X}_{i,t} : i \in \mathcal{I} \text{ and } t \leq s\}$. Further, we assume that this conditional dependency is k -Markovian in nature, that is

$$\mathbb{P}(\mathcal{Y}_s | \mathcal{X}_{i,t} : i \in \mathcal{I} \text{ and } t \leq s) = \mathbb{P}(\mathcal{Y}_s | \mathcal{X}_{i,t} : i \in \mathcal{I} \text{ and } s - k \leq t \leq s).$$

For this work, we chose k to be 28 days, meaning distributions of precipitation at a location only depend on the values of the model fields from the last 28 days, which is a reasonable assumption from a meteorological point of view.

Since this conditional joint density is hard to estimate in one step, we decompose our approach into two sub-parts using Sklar’s theorem (Sklar, 1959) on copulas, which are multivariate cumulative distribution functions (CDFs) for which the marginal probability distribution of each variable is uniform on the interval $[0, 1]$.

Sklar’s theorem uses copula densities on $[0, 1]^n$ to link the marginal density functions from n dimensions together, thereby forming a joint probability density function (PDFs) for random variables in \mathbb{R}^n . But we notice that the observed rainfall $y_{i,t}$ is a non-negative random variable with a significantly high probability of zero, as a rainfall value for a dry

day translates to a zero-valued observation, making our observations a case of *zero-inflated data*, in which zero-valued data points are over-represented.

Hence, to apply Sklar's theorem for modelling the joint distribution of rainfall, we consider our observations \mathbf{Y}_s being censored observations of a latent variable $\mathbf{Y}_s^* \in \mathbb{R}^n$ where

$$y_{i,s} = \begin{cases} y_{i,s}^*, & \text{if } y_{i,s}^* > 0 \\ 0, & \text{if } y_{i,s}^* \leq 0 \end{cases}.$$

We assume the conditional marginal densities of the latent variables \mathbf{Y}_s^* are $f_{i,s}^*(y_{i,s}^*|\mathcal{H}_s)$, $i \in \mathcal{I}$, $s \in \mathcal{T}$ (with distributions $F_{i,s}^*(y_{i,s}^*|\mathcal{H}_s)$) at location $i \in \mathcal{I}$ and model the dependence between them with a unique copula

$$\mathbf{c}^* (F_{1,s}^*(y_{1,s}^*|\mathcal{H}_s), \dots, F_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) | \mathcal{H}_s, \mathcal{D})$$

given \mathcal{H}_s and \mathcal{D} , the latter being a $n \times n$ matrix of spatial information for n locations in \mathcal{I} . Hence Sklar's theorem (Sklar, 1959) provides us with a joint density, denoted $\pi(\mathbf{Y}_s^*|\mathcal{H}_s)$, of the form

$$\pi^*(\mathbf{Y}_s^*|\mathcal{H}_s) = f_{1,s}^*(y_{1,s}^*|\mathcal{H}_s) \cdots f_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) \cdot \mathbf{c}^* (F_{1,s}^*(y_{1,s}^*|\mathcal{H}_s), \dots, F_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) | \mathcal{H}_s, \mathcal{D}).$$

We consider $f_{i,s}$ and $F_{i,s}$ are respectively the conditional marginal density and conditional marginal cumulative distribution function of rainfall at location i and time s given \mathcal{H}_s , which corresponds to the probability density function (with respect to the Lebesgue $+\delta_0$ measure)

$$f_{i,s}(y_{i,s}|\mathcal{H}_s) = \begin{cases} f_{i,s}^*(y_{i,s}|\mathcal{H}_s), & \text{if } y_{i,s} > 0 \\ F_{i,s}^*(0|\mathcal{H}_s), & \text{if } y_{i,s} = 0 \end{cases}$$

and the distribution function

$$F_{i,s}(y_{i,s}|\mathcal{H}_s) = \begin{cases} F_{i,s}^*(y_{i,s}|\mathcal{H}_s), & \text{if } y_{i,s} > 0 \\ p_{i,s} = F_{i,s}^*(0|\mathcal{H}_s), & \text{if } y_{i,s} = 0 \end{cases}.$$

In order to model the joint distribution of rainfall, our approach at first involves directly modelling the marginal densities $f_{i,s}$ (and correspondingly, distributions $F_{i,s}$) by leveraging established distributions of zero-inflated data. Examples of such distributions include a mixture model with mass at $\{0\}$ and a continuous parametric density for $y_{i,t} > 0$ (Das, 1955; Shimizu, 1993). To condition the marginal densities $f_{i,s}$ on \mathcal{H}_s , we consider the parameters of this distribution family to be functions of \mathcal{H}_s in our novel approach termed the joint generalised neural model (JGNM). This model transforms \mathcal{H}_s into informative summaries via a function parameterised by a neural network architecture. These summaries are then employed as predictors within a joint generalised linear model framework (Nelder and Lee, 1991) to jointly model the different parameters of the marginal densities, hence the name. The initial transformation of predictors into summaries is designed to extract non-linear trends from the data. More specifically, the temporal and spatial dependency structures




Figure 1: **Diagram of our Cens-JGNN model:** Our method takes weather variables and distances between locations as inputs with the goal of downscaling samples to higher resolution target rainfall. In part 1), we predict parameters for marginal distributions. In part 2), we model the dependence with a censored latent copula. Finally, in part 3), we produce rainfall samples by applying the inverse distributions to latent copula samples.

are captured through distinct neural network architectures, employing convolutional long short-term memory network (Shi et al., 2015) modules. For each of the locations in \mathcal{I} and times in \mathcal{T} , the JGNN is trained to output parameters which minimize the negative log-likelihood of rainfall under the parametric family model for $f_{i,s}(y_{i,s}|\mathcal{H}_s)$. The detailed description of JGNN and the inference process are elaborated in Section 3.

In the second part, to model the spatial dependency structure we assume a Gaussian copula for the latent continuous variables \mathcal{Y}_s^* . To encapsulate the spatial dependency structure among the n locations, we draw upon the literature on Gaussian processes (Rasmussen et al., 2006). The correlation matrix Σ of the Gaussian copula is determined by a covariance function k , which offers a scalar measure of similarity between the spatial information of two points. This measure is often referred to as the kernel of the Gaussian process. The correlation matrix is generated by pairwise evaluation of the kernel on all points. For the sake of simplicity, we assume that the spatial dependency structure is time-independent,

and hence Σ does not depend on \mathcal{H}_s . The specifics of the parametric form and the inference of the kernel parameter will be discussed in Section 4. It is important to note that the computation of the maximum likelihood estimator of the kernel parameters is infeasible as we do not have knowledge of $f_{i,s}^*(y_{i,s}^*)$ and $F_{i,s}^*(y_{i,s}^*)$ for $y_{i,s}^* \leq 0, \forall i \in \mathcal{I}$. As a result, we propose an innovative estimation methodology based on the minimum Scoring Rules estimator (Dawid et al., 2016; Pacchiardi and Dutta, 2022b; Alquier et al., 2022), which requires only simulations from the censored latent Gaussian copula. We provide a diagram of our model, termed Cens-JGNM, in Figure 1.

3. Modelling Marginals with the Joint Generalised Neural Model

We will now describe the considered parametric model for the conditional marginal densities $f_{i,s}(y_{i,s}|\mathcal{H}_s)$ and their corresponding distributions $F_{i,s}(y_{i,s}|\mathcal{H}_s)$ in Section 3.1. Later in Section 3.2, we describe how to train parameters of these densities with details of our model's neural architecture given in Section 3.3.

3.1 Joint Generalised Neural Model

As our goal is to retain explainability in our modelling approach, for marginal densities, we opt for a parametric form. Additionally, this saves computational costs in Section 4, where repeated exact evaluations of the distributions are necessary. Due to the zero-inflated nature of rainfall data, we consider the following parametric form for the marginal density

$$\mathcal{P}(y_{i,s}; p, \mu_{i,s}, \phi_{i,s}) = [1 - p_{i,s}] \cdot \delta_0(y_{i,s}) + p_{i,s} \cdot \mathcal{F}(y_{i,s}; \mu_{i,s}, \phi_{i,s}),$$

for $y_{i,s} \in \mathbb{R}_{\geq 0}, i \in \mathcal{I}, s \in \mathcal{T}$, where $\delta_0(\cdot)$ is the Dirac mass measure at 0, $p_{i,s} \in [0, 1]$ is the probability of positive rainfall and with $\mu_{i,s} > 0$ and $\phi_{i,s} > 0$ symbolising the mean and dispersion parameters¹ of \mathcal{F} which is a continuous density defined on $\mathbb{R}_{>0}$. Existing studies (Das, 1955; Fealy and Sweeney, 2007; Little et al., 2009; Ailliot et al., 2015; Holsclaw et al., 2017; Bertolacci et al., 2019; Xie et al., 2023) argue in favour of a gamma distribution as the choice of \mathcal{F} , following which we will consider the zero-gamma mixture in the rest of this work with density given by

$$\mathcal{P}(y_{i,s}; p, \mu_{i,s}, \phi_{i,s}) = [1 - p_{i,s}] \cdot \delta_0(y_{i,s}) + p_{i,s} \cdot \left(\frac{y_{i,s}}{\phi_{i,s} \cdot \mu_{i,s}} \right)^{1/\phi_{i,s}} \frac{1}{y_{i,s}} \exp\left(-\frac{y_{i,s}}{\phi_{i,s} \cdot \mu_{i,s}}\right) \frac{1}{\Gamma(1/\phi_{i,s})},$$

for $y_{i,s} \in \mathbb{R}_{\geq 0}, i \in \mathcal{I}, s \in \mathcal{T}$. We further assume that the marginal distributions are influenced by the conditioning predictor variables (\mathcal{H}_s) exclusively through their parameters.

For any given $y_{i,s}$ and a relevant subset of predictors $\mathbf{x}_{i,s} \subseteq \mathcal{H}_s$, following the literature on generalised linear models (GLMs) (Dunn and Smyth, 2018) and joint generalised linear model (JGLMs) (Nelder and Lee, 1991), we can model the parameters $(p_{i,s}, \mu_{i,s}, \phi_{i,s})$ specific to location i and time t as,

$$\begin{aligned} g_p(p_{i,s}) &= \alpha_0 + \langle \boldsymbol{\alpha}, \mathbf{x}_{i,s} \rangle \\ g_\mu(\mu_{i,s}) &= \beta_0 + \langle \boldsymbol{\beta}, \mathbf{x}_{i,s} \rangle \\ g_\phi(\phi_{i,s}) &= \gamma_0 + \langle \boldsymbol{\gamma}, \mathbf{x}_{i,s} \rangle \end{aligned} \tag{1}$$

1. Note here that $\phi_{i,s}$ refers to the dispersion parameter while ϕ without subscripts refers to the Gaussian PDF.

where $(\alpha_0, \boldsymbol{\alpha})$, $(\beta_0, \boldsymbol{\beta})$ and $(\gamma_0, \boldsymbol{\gamma})$ are regression coefficients and $g_{p_{i,s}}(p_{i,s}) = \log(p_{i,s}/(1 - p_{i,s}))$, $p_{i,s} \in (0, 1)$, $g_{\mu_{i,s}}(\mu_{i,s}) = \log(\mu_{i,s})$, $g_{\phi_{i,s}}(\phi_{i,s}) = \log(\phi_{i,s})$ are the link functions. In practice, since we do not require reversibility of our link function, we are free to choose an inverse link function that has an appropriate range, but provides a more stable set of gradients during training. Table 5 details the specific inverse link function used in our work.

While a JGLM would allow us to have location and time-dependent models for marginal densities $f_i(y_{i,s}|\mathcal{H}_s)$, we can still seek improvement in the complexity of the model. Notably, the relationship between predictors and density parameters is linear in the coefficients. As such, non-linear relationships in the modelling of parameters will not be taken into account by the marginals. To remedy this, a further improvement of GLMs can be considered by incorporating neural networks (Goodfellow et al., 2016).

Similar to the DeepGLM model introduced in Tran et al. (2020), we refine the predictors $\mathbf{x}_{i,s} \subseteq \mathcal{H}_s$ by training a neural network \mathcal{N}_ψ parameterised by ψ to obtain refined features $\mathcal{N}_\psi(\mathbf{x}_{i,s})$. These refined features are used in the place of $\mathbf{x}_{i,s}$ as predictor variables in the Equation (1) to arrive at our joint generalised neural model,

$$\begin{aligned} g_p(p_{i,s}) &= \alpha_0 + \langle \boldsymbol{\alpha}, \mathcal{N}_\psi(\mathbf{x}_{i,s}) \rangle \\ g_\mu(\mu_{i,s}) &= \beta_0 + \langle \boldsymbol{\beta}_i, \mathcal{N}_\psi(\mathbf{x}_{i,s}) \rangle \\ g_\phi(\phi_{i,s}) &= \gamma_0 + \langle \boldsymbol{\gamma}, \mathcal{N}_\psi(\mathbf{x}_{i,s}) \rangle \end{aligned}$$

where $(\alpha_0, \boldsymbol{\alpha})$, $(\beta_0, \boldsymbol{\beta})$ and $(\gamma_0, \boldsymbol{\gamma})$ are regression coefficients in \mathbb{R} and $g_p(\cdot)$, $g_\mu(\cdot)$, $g_\phi(\cdot)$ represent the link functions.

Thus, our JGNM model can be seen as a distributional learner capable of outputting multiple parameters given initial weather predictors by capturing linear as well as non-linear relationships present in the data. Unlike models that output a field of predictions over fixed points (e.g. Harris et al. (2022); Vosper et al. (2023)), our model only needs predictors to output parameterised densities. Thus, given predictors $\mathbf{x}_{i,s}$ for a new location i expected to follow the same precipitation patterns as in \mathcal{I} , our model can be applied to such locations outside the training set, thereby enabling, for example, principled interpolation. Indeed, one only needs to input initial predictors $\mathbf{x}_{i,s}$ into \mathcal{N}_ψ in order to have the JGNM parameterise a distribution; no further distinction is made between locations given the input, meaning additional locations can be included into our model without needing to retrain it.

3.2 Inference for the JGNM

To ensure our model aligns with the selected distributional assumptions, we learn the aggregated parameters $(\psi, \alpha_0, \boldsymbol{\alpha}, \beta_0, \boldsymbol{\beta}, \gamma_0, \boldsymbol{\gamma})$ using the likelihood specific to the assumed marginal model in the loss function of the JGNM.

For the gamma mixture, the JGNM loss function contains two segments. The first segment pertains to $p_{i,s}$ and adopts the form of a logistic loss. The second segment is associated with the parameters $\mu_{i,s}, \phi_{i,s}$ and is represented by the negative log-likelihood (NLL) of the corresponding gamma density. In instances of zero rainfall, only the logistic component is considered, while in the event of rainfall, both components are evaluated.

The component in the JGNM loss function accounting for rainfall occurrence in terms of $p_{i,s}$ is the logistic loss, also known as binary cross entropy, between occurrence and $p_{i,s}$,

written as

$$\mathcal{L}(p_{i,s}|y_{i,s}) = -\mathbb{I}\{y_{i,s} = 0\} \log(p_{i,s}) - (1 - \mathbb{I}\{y_{i,s} = 0\}) \log(1 - p_{i,s}).$$

Under a GLM parameterisation the Gamma density for positive rain $y_{i,s} > 0$ is written as

$$\begin{aligned} f_i(y_{i,s}|\mathcal{H}_s) &= \mathcal{P}(y_{i,s}; p_{i,s}, \mu_{i,s}, \phi_{i,s}) \\ &= \left(\frac{y_{i,s}}{\phi_{i,s}\mu_{i,s}}\right)^{1/\phi_{i,s}} \frac{1}{y_{i,s}} \exp\left(-\frac{y_{i,s}}{\phi_{i,s}\mu_{i,s}}\right) \frac{1}{\Gamma(1/\phi_{i,s})} \end{aligned}$$

leading to a negative log-likelihood corresponding to the loss for positive rain of the form

$$\mathcal{L}(\mu_{i,s}, \phi_{i,s}|y_{i,s}) = \left(\frac{1}{\phi_{i,s}} - 1\right) \log(y_{i,s}) - \frac{\log(\phi_{i,s}\mu_{i,s})}{\phi_{i,s}} - \frac{y_{i,s}}{\mu_{i,s}\phi_{i,s}} - \log(\Gamma(\phi_{i,s}^{-1})).$$

3.3 Network Architecture

The neural network component $[\mathcal{N}_\psi]$ of the JGNM employs a unique architecture to capture the spatio-temporal properties of the data. Our proposed model contains a convolutional long short-term memory network (ConvLSTM) (Shi et al., 2015) — a neural module that combines convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) — to extract both spatial and temporal trends effectively. ConvLSTM leverages the strengths of LSTM models, which are recursive models appropriate for time-series data, and convolutional networks, a notable architecture for 2D/3D image data. A comprehensive description of the architecture can be found in Appendix A.

In addition to the standard ConvLSTM, our architecture is novel in its integration of several key modules that differentiate it from conventional approaches. Notably, the Time Distributed 2D Convolution Layer (TD2L) ensures uniform feature extraction across time steps, while the Time Distributed Temporal Downscaling layer (TDTD) employs a convolutional self-attention mechanism to dynamically aggregate temporal information. Furthermore, the Time Distributed Mean Function Layer, composed of dedicated single-layer perceptrons, provides tailored inverse link functions for each parameter of the predictive distribution. Instead of employing simple temporal aggregation, we introduce a Convolutional Self-Attention layer to facilitate the transition within our model from 6-hourly weather field data to a 24-hourly representation.

Overall, our network comprises several key components with specific sizes and layer counts. The initial TD2L uses one 2D convolution to transform $16 \times 16 \times 6$ input patches into $16 \times 16 \times 64$ feature maps. This is followed by two bi-directional ConvLSTM layers with skip connections to robustly capture spatio-temporal dependencies. The TDTD layer then groups every four consecutive time steps, employing a CSA module that uses 4 attention heads (each with a hidden dimension of 64) to compute dynamic, weighted feature aggregations. This multi-head CSA module enhances the model’s ability to integrate information across time before the final downscaling. Finally, the Time Distributed Mean Function Layer uses three perceptrons to produce the parameters of our predictive distribution.

4. Censored Latent Gaussian Copula

In this section, we centre our attention on the dependency structure model described in Section 4.1 and the associated challenges in its estimation, as discussed in Section 4.2. Our proposed solution is a novel estimation approach based on minimum scoring rules detailed in Section 4.3, which we validate on simulated data in Section 4.4.

For the latent variables \mathbf{y}_s^* , we assume that their dependence can be modelled by a Gaussian copula \mathbf{c} , with density

$$\begin{aligned} & \mathbf{c}^* (F_{1,s}^* (y_{1,s}^* | \mathcal{H}_s), \dots, F_{n,s}^* (y_{n,s}^* | \mathcal{H}_s) | \mathcal{H}_s, \mathcal{D}) \\ &= \mathbf{c}(\Phi(z_{1,s}^*), \dots, \Phi(z_{n,s}^*) | \Sigma(\mathcal{D} | \theta)) = \frac{\phi_n(z_{1,s}^*, \dots, z_{n,s}^* | \mathbf{0}, \Sigma(\mathcal{D} | \theta))}{\prod_{i=1}^n \phi(z_{i,s}^*)} \end{aligned}$$

where $z_{i,s}^* = \Phi^{-1} (F_{i,s}^* (y_{i,s}^* | \mathcal{H}_s))$, ϕ and Φ are the PDF and CDF of univariate Gaussian distributions, and ϕ_n a Gaussian multivariate density with mean $\mathbf{0}$ and correlation matrix $\Sigma(\mathcal{D} | \theta)$ which depends on a matrix of spatial information \mathcal{D} through a parameter θ . We construct \mathcal{D} as the Euclidean distance based on the latitude and longitude of locations, the details of which are provided in Appendix B where we also investigate and decide against using topographical information for \mathcal{D} . We assume the dependence structure imposed by this correlation to be constant in time, hence dropping the dependence on \mathcal{H}_s in the construction of Σ . This is a tradeoff of simplicity for modelling purposes and sophistication for predictive prowess. In this way, we model the joint *relative* intensity as a function of space, while the exact amounts also depend on temporal information through the marginals. Due to the Gaussianity assumption on \mathbf{c} , the only object to infer is the correlation matrix which incorporates the spatial dependence into our model.

4.1 Modelling of Spatial Dependency

This reliance on \mathcal{D} when constructing the latent Gaussian copula \mathbf{c} induces the spatial dependence on latent variables \mathbf{y}_s^* and hence on the observations \mathbf{y}_s . To construct the correlation, we reference the existing literature on Gaussian processes (Rasmussen et al., 2006) and apply a Gaussian process kernel as an element-wise transformation of the initial matrix \mathcal{D} , which contains spatial information. This transformation ensures the resulting matrix's positive definiteness and thus its validity as a covariance of a Gaussian distribution. The choice of kernel $k(\mathcal{D} | \theta)$ specifies this transformation and typically depends on a few parameters, here denoted by θ . We use the Matérn kernel (Matérn, 2013), a common choice in spatial statistics (Gerber and Nychka, 2021; Richards et al., 2024), which provides us with a correlation matrix $\Sigma(\mathcal{D} | \theta)$ of the following form:

$$\Sigma_{(i,j)}(\mathcal{D}_{(i,j)} | \theta) = k(\mathcal{D}_{(i,j)} | \theta) = \frac{2^{1-\nu}}{\Gamma(\nu)} \cdot \left(\sqrt{2\nu} \frac{\mathcal{D}_{(i,j)}}{\theta} \right)^\nu \cdot K_\nu \left(\sqrt{2\nu} \frac{\mathcal{D}_{(i,j)}}{\theta} \right)$$

where Γ is the gamma function and K_ν is the modified Bessel function of the second kind, noting that the expression depends on the unknown lengthscale $\theta (\in \mathbb{R}_{>0})$ and additional parameter ν usually chosen by the user. In this work, we would consider ν to be fixed at 0.5 as for values of $\nu = n + 1/2, n \in \mathbb{N}$ the Matérn kernel has an equivalent simpler expression

(as an exponential multiplied by a polynomial of order n), granting us a computational speedup (Guttorp and Gneiting, 2006) (see details in Appendix D).

A convenient consequence of such a construction for the correlation matrix is that it becomes straightforward to include additional locations (if we expect them to follow the same pattern as our training set) in our approach even after fitting it to data. Indeed, as long as one has relevant spatial information about a new location, the initial distance matrix \mathcal{D} can be expanded to include that new location by applying the learned kernel to the distances, including it in the dependence structure.

4.2 Infeasibility of Maximum Likelihood Estimation

Traditionally, inference for a Gaussian copula is done using maximum likelihood estimation (MLE) on its density. However, our Gaussian copula is latent, meaning we do not have access to direct realisations $z_{i,s}^* = \Phi^{-1}\left(F_{i,s}^*(y_{i,s}^*|\mathcal{H}_s)\right)$ assumed to come from it. Instead, we examine the joint density of observed data conditional on the history $\pi(\mathcal{Y}_s|\mathcal{H}_s)$. It can be obtained by considering the joint density of the latent variables $\pi^*(\mathcal{Y}_s^*|\mathcal{H}_s)$ and marginalising over censored locations denoted as $\mathcal{C}_s = \{i \in \mathcal{I} : y_{i,s} = 0\}$. The marginalisation is done up to the censoring thresholds of 0, to give the marginalised mass under the latent joint density $\pi^*(\cdot|\mathcal{H}_s)$ to censored events under $\pi(\cdot|\mathcal{H}_s)$ in the form of point mass. This is formalised as:

$$\begin{aligned} \pi(\mathcal{Y}_s|\mathcal{H}_s) &= \int_{-\infty}^0 \dots \int_{-\infty}^0 \pi^*(\mathcal{Y}_s^*|\mathcal{H}_s) \prod_{i \in \mathcal{C}_s} dy_{i,s}^* \\ &= \int_{-\infty}^0 \dots \int_{-\infty}^0 f_{1,s}^*(y_{1,s}^*|\mathcal{H}_s) \cdot \dots \cdot f_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) \\ &\quad \cdot \mathbf{c}^*(F_{1,s}^*(y_{1,s}^*|\mathcal{H}_s), \dots, F_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) | \mathcal{D}, \theta) \prod_{i \in \mathcal{C}_s} dy_{i,s}^* \\ &= \int_{-\infty}^0 \dots \int_{-\infty}^0 f_{1,s}^*(y_{1,s}^*|\mathcal{H}_s) \cdot \dots \cdot f_{n,s}^*(y_{n,s}^*|\mathcal{H}_s) \\ &\quad \cdot \frac{\phi_n(z_{1,s}^*, \dots, z_{n,s}^* | \mathbf{0}, \Sigma(\mathcal{D}|\theta))}{\prod_{i=1}^n \phi(z_{i,s}^*)} \prod_{i \in \mathcal{C}_s} dy_{i,s}^*. \end{aligned}$$

While we know the form of the copula by assumption, performing MLE on this expression is in-feasible due to the unavailability of marginals $f_{i,s}^*, F_{i,s}^*$ over $(-\infty, 0]$ which are needed for the integration over censored values. Thus, we cannot rely on MLE for inference of the kernel parameter θ .

To circumvent the problematic intractable likelihood expression, we develop a methodology using minimum scoring rules estimation (Dawid et al., 2016) for the inference of parameters of Gaussian copulas, suitable to our case of censored latent Gaussian copulas.

4.3 Minimum Scoring Rule Censored Gaussian Copula

Introduced in Dawid et al. (2016); Pacchiardi and Dutta (2022a), the idea of minimum scoring rule estimation is to simulate realisations from the model for a chosen parameter

$\theta \in \Theta$ and select the best such parameter by comparing the simulated data against observations. This comparison is done using scoring rules (Gneiting and Raftery, 2007) which under suitable conditions define a divergence measure between two distributions and can be generalised to compare two data sets. As defined in Gneiting and Raftery (2007), a scoring rule $S(\mathcal{P}_\theta, \mathbf{z})$ is a function between a distribution \mathcal{P}_θ and observed data \mathbf{z} as a realisation of a random variable $\mathbf{Z} \sim \mathcal{P}^*$. Hence a minimum scoring rule estimate becomes,

$$\hat{\theta} = \arg \min_{\theta \in \Theta} S(\mathcal{P}_\theta, \mathbf{z}).$$

If the scoring rule is strictly proper (as defined in Appendix C), this can be shown as asymptotically (over repeated samples \mathbf{z}) equivalent to minimizing a statistical divergence between \mathcal{P}_θ and \mathcal{P}^* . Considering the negative log probability density function of \mathcal{P}_θ evaluated at the observation as a score function, we can arrive at the maximum likelihood estimation scheme and correspondingly to the Kullback–Leibler divergence. We provide more details regarding these connections in Appendix C. Due to the intractability of the probability density function of the latent Gaussian copula, we choose the energy score,

$$S_E^{(\beta)}(P, \mathbf{z}) = 2 \cdot \mathbb{E} \left[\|\mathbf{Z} - \mathbf{z}\|_2^\beta \right] - \mathbb{E} \left[\|\mathbf{Z} - \mathbf{Z}'\|_2^\beta \right], \quad \mathbf{Z} \perp \mathbf{Z}' \sim P,$$

which is a strictly proper scoring rule for $\beta \in (0, 2)$. In our study, we fix $\beta = 0.5$ for better optimisation stability based on preliminary experiments, which we will assume henceforth for all applications of the energy score. If we can draw m identical and independent samples $\mathbf{z}' = (\mathbf{z}'_1, \dots, \mathbf{z}'_m)$ from the distribution \mathcal{P}_θ , then an unbiased estimate of the energy score $S_E^{(\beta)}(\mathcal{P}_\theta, \mathbf{z})$ can be constructed as follows:

$$\hat{S}_E^{(\beta)}(\mathbf{z}', \mathbf{z}) = \frac{2}{m} \sum_{j=1}^m \|\mathbf{z}'_j - \mathbf{z}\|_2^\beta - \frac{1}{m(m-1)} \sum_{\substack{k=1 \\ k \neq j}}^m \|\mathbf{z}'_j - \mathbf{z}'_k\|_2^\beta.$$

The inference of the parameter for the censored latent Gaussian copula model depends on the crucial observation that we can simulate from this model for any given parameter θ . As a Gaussian copula is associated to a Gaussian with a given correlation matrix Σ (Nelsen, 2006), we can infer the correlation matrix directly on the Gaussian scale by transforming the observed data to the Gaussian scale as $z_{i,s} = \Phi^{-1}(F_{i,s}(y_{i,s} | \mathcal{H}_s)) \in \mathbb{R}$. For a censored Gaussian copula, we correspondingly deal with a censored Gaussian density. To compute the minimum scoring rule estimate of θ , we define simulations $\mathbf{z}'_s = (z'_{1,s}, \dots, z'_{n,s})$ from the censored Gaussian density (corresponding to the censored Gaussian copula) as

$$z'_{i,s} = \begin{cases} z_{i,s}^* & \text{if } z_{i,s}^* > d_{i,s} \\ d_{i,s} & \text{if } z_{i,s}^* \leq d_{i,s} \end{cases} \quad (2)$$

where $\mathbf{z}'_s = (z'_{1,s}, \dots, z'_{n,s})$ is a simulation from the uncensored Gaussian density $\mathcal{N}(\cdot | \mathbf{0}, \Sigma(\mathcal{D}|\theta))$, $d_{i,s} = \Phi^{-1}(1 - p_{i,s})$ and $p_{i,s}$ corresponds to the mass at 0 for the marginal density model $F_{i,s}(\cdot | \mathcal{H}_s)$. Further, we also transform observed rainfall $(y_{i,s} : 1 \leq i \leq n)$ to $\mathbf{z}_s = (z_{i,s} : 1 \leq i \leq n)$ which is the corresponding output of the (assumedly true) censored Gaussian. We do so by applying the inferred marginal distribution $F_{i,s}(\cdot | \mathcal{H}_s)$ to observation $y_{i,s}$, getting




Figure 2: **Simulation study:** We fit a regular and a censored copula for 1000 replicates with our method, repeating this for increasing sample sizes. (a) The parameter estimation error distributions for both cases with the associated MSE, with the dotted lines representing a Gaussian approximation of the error density. (b) The L_1 norm compares the true and estimated matrix obtained with the fitted parameter $\hat{\theta}$. Both the censored and uncensored copulas show a decreasing error with the observed data size, occurring faster for the latter.

$z_{i,s} = \Phi^{-1}(F_{i,s}(y_{i,s}|\mathcal{H}_s))$ at each location i and time s . Comparing the transformed observations \mathbf{z}_s with m samples from the censored Gaussian with correlation parameterised by θ , denoted by $\mathbf{Z}'_{\theta,s} = \left\{ \mathbf{z}'_s^{(j)} = (z'_{1,s}, \dots, z'_{n,s})^{(j)} \right\}_{j=1}^m$, for each $s \in \mathcal{S}$, we can compute the minimum scoring rule estimate as follows:

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \sum_{s \in \mathcal{S}} \hat{S}_E^{(\beta)}(\mathbf{z}'_{\theta,s}, \mathbf{z}_s).$$

4.4 Inference on simulated data

To empirically validate our estimation procedure, we consider a simulation study with known parameter values. The experimental setup is as follows. For a given observed sample size of $\mathcal{T} = 250, 500, 750, 1000$ we sample from a spatial Gaussian copula with a correlation matrix Σ^* parameterised using $\theta^* = 35$ as described in Section 4.1, and from a censored Gaussian copula with identical correlation where the censoring levels are $\{1 - p_i\}_{i=1}^{|\mathcal{I}|} \sim U[0.5, 0.95]$. We fix $|\mathcal{I}| = 3$ since under a correctly specified model the correlation kernel is applied entry-wise, hence rendering all dimension pairs informative about θ , meaning inference can be conducted with any subset of locations. We optimise θ using our Minimum Scoring Rule approach of Section 4.3 with the gradient-based Adam (Kingma and Ba, 2014) optimiser to convergence to obtain our parameter estimate $\hat{\theta}$, in turn getting $\hat{\Sigma}$. The size of simulations for the Scoring Rule is always set to the observation sample size. We repeat this experiment 1000 times with different optimisation initialisations in $U[30, 40]$ to imitate our initialisation in practice, see Appendix D.

In Figure 2a, we show the distribution of estimation errors for θ with sample size as well as the mean squared error, in blue and red respectively for the censored and uncensored Gaussian copulas. Our minimal SR inference approach recovers the true parameter for both copula types, with the estimation error shrinking with sample size and approximately following a Gaussian density (shown as dotted lines fitted on the errors) centred at the true parameter value. In Figure 2b, we show the L_1 norm (see Appendix E.1) of correlation matrix differences $\Sigma_{\hat{\theta}} - \Sigma_{\theta^*}$, depicting a similar pattern, where the estimated correlation matrix matches the true matrix better with increasing sample sizes.

In the uncensored case, the error variance decreases more rapidly than in the censored case. This reflects the loss of information due to censoring. Practically, instead of observing the full range of copula samples $[0, 1]^n$, we only observe the range $[1-p_{1,s}, 1] \times \dots \times [1-p_{n,s}, 1]$, making only a fraction of the samples fully observed, effectively diminishing our sample size. While the censored data still carries information about the correlation structure, there is a loss of information, which explains the slower rate of decreasing error in Figure 2b.

5. Real Data Scenario

Finally, we apply our approach to the data set presented in Section 2. We conduct a comprehensive evaluation by comparing our proposed method with other benchmark probabilistic downscaling methods using standard diagnostics of probabilistic prediction [Appendix E] in Section 5.1, in addition to a deterministic precipitation prediction provided by IFS data. We then assess the spatial coherence of the generated samples from all these methods in Section 5.2. Finally, we investigate the effect of varying lengths of training data on the performance of our approach in Section 5.3.

5.1 Comparison with benchmark models

We begin by comparing our method (denoted Cens-JGNM) to two competing benchmark methods: the variational auto-encoder generative adversarial network (VAE-GAN) (Harris et al., 2022) and the convolutional conditional neural process (ConvCNP) (Vaughan et al., 2022). Both these methods as well as our own were fitted to the grided data detailed in Section 2.1 spanning a 20-year period, divided into a training set from 1979 to 1993 and a validation set from 1993 to 1999. We then compare their forecasting capabilities on a held-out test set of 20 years, from 1999 to 2019.

Remark: It is worth reiterating that our data does not include precipitation weather forecasts as a predictor, unlike the above-mentioned studies of Harris et al. (2022); Vaughan et al. (2022). Consequently, the performance of the benchmark methods obtained in their respective articles cannot be directly compared with the performances presented in this study (of the benchmarks as well as our approach). The inclusion of precipitation NWPMs predictions would improve a model’s forecasting capabilities for rainfall.

We begin by looking at the Area Under the Curve (AUC) [Appendix E.2] resulting from a calibration task, as shown in Figure 3a, with corresponding Receiver Operating Characteristic (ROCs) curves reported in Appendix F. The AUC achieved by the Cens-JGNM is slightly higher than that of the ConvCNP, and both remain constant across different precipitation amounts. The VAE-GAN’s AUC is generally lower than the two




Figure 3: **Calibration diagnostics:** We show in (a) AUCs [Appendix E.2], (b) Rank histograms [Appendix E.3] of the rank of precipitation observations against simulated samples and (c) ESFs [Appendix E.4] for our Cens-JGNM approach compared to ConvCNP and VAE-GAN benchmarks. Figures 3a and 3b indicate the appropriate treatment of marginal densities by the Cens-JGNM and ConvCNP both employing explicit Gamma densities as opposed to the implicit approach of the VAE-GAN. Figure 3c reveals that the Cens-JGNM obtains the best calibration, closely followed by the ConvCNP, with the VAE-GAN incorrectly capturing 0 values and underpredicting higher amounts.

other methods and deteriorates for higher precipitation amounts, showing that the two aforementioned models are better suited for modelling high rainfall events.

In Figure 3b, we show rank histograms [Appendix E.3] to investigate calibration. The Cens-JGNM and ConvCNP have almost identical performances, displaying a close fit to the ideal rank, which we attribute to their similar form of marginal densities with the adoption of a Gamma mixture model. The slight misalignment of ranks on the far right for the two aforementioned approaches suggests that the parametric form is not entirely accurate in the tail, and could benefit from a more appropriate treatment of extreme values. On the other hand, the VAE-GAN displays a poor performance characteristic of under-dispersion with ranks over-represented at 0 and 1, arguing in favour of a parametric approach for marginal densities.

The Empirical Survival Functions (ESFs) [Appendix E.4] of each approach are studied in Figure 3c. The Cens-JGNM has the closest fit to the empirical frequency, suggesting it has the best calibration out of the three models. As this diagnostic is aggregated over space, we attribute ConvCNP’s overestimation of realised rainfall to the lack of explicit dependence modelling leading to miss-calibrated joint forecasts. The under-dispersion and non-explicit modelling of zero values of the VAE-GAN is exemplified with an incorrect intercept at 0 *mm* and a lower frequency for higher amounts.

5.2 Spatial coherence

To assess the spatial fidelity of each model, we compute multiple metrics detailed in Appendix E. Qualitatively, we also show samples from each model for three different days in Figure 4. Visually, the ConvCNP displays poor spatial dependence as rainfall values vary sharply between neighbouring locations. The VAEGAN and Cens-JGNM samples display adequate spatial smoothness. Notably, all models have a large variability in their samples,





Figure 4: **Rainfall forecasts:** For a given target day at the top of each column, we show row by row samples from the Cens-JGNM, the VAE-GAN and the ConvCNP models. Each model’s quadrant top left image is the median forecast, with the remaining three figures being random samples for that day. The Cens-JGNM and VAEGAN samples look spatially consistent while the ConvCNP samples are disparate across space.



(a) Covariance of rainfall against distance




Figure 5: **Spatial and temporal coherence:** We show covariance of rainfall against distance (a), for observed data and all three models. We show a spatio-temporal energy score across regions of London (b), Glasgow (c) and Dublin (d). As evidenced by plot (a), our Cens-JGNM is the most effective approach at capturing spatial dependence while the VAE-GAN and ConvCNP capture minimal dependence only. Plots (b,c,d) show the stronger spatio-temporal predictive performance of our method over benchmarks.

rarely matching the target rain exactly, possibly indicating the stochastic nature of the rainfall downscaling task.

For quantitative results, we begin by computing the average covariance across distances in simulated samples from each model and observed data, shown in Figure 5a. For a given distance (on the x-axis), for all the pairs of locations i, j exactly that far apart, we compute the sample covariance (on the y-axis) between the values at those two locations over all days in \mathcal{T} . We show the average covariance over all such pairs i, j as a line with confidence bands as the standard deviation over all these pairs. Note that for higher distances, the number of locations that distance apart greatly decreases, thus leading to more variable covariance estimates despite the confidence bands shrinking.




Figure 6: **Spectral Distribution:** Average ratio of spectral energies across wavenumbers between predictions for each model and the true observed rainfall values. The high resolution deterministic benchmark IFS demonstrates the best overall performance, while Cens-JGNM shows superior results among neural network-based approaches and is similar to the IFS, particularly at higher wavenumbers.

The Cens-JGNM overall matches the observed covariance the closest of all models, being aligned to the observed covariance for all distances. The Cens-JGNM and ConvCNP both slightly overestimate the covariance for large distances, while the VAEGAN underestimates the covariance significantly across all distances. The standard deviation of the Cens-JGNM is lower than for observed data, meaning that our model does not greatly differentiate between different pairs and only cares about the distance between them. This is expected since distance is the only input into our copula covariance matrix which is the main factor behind the covariance.

Next, in Figures 5b,5c, and 5d, we compute the aggregate rainfall across the set of locations surrounding London, Glasgow and Dublin respectively, over multiple days as explained in Appendix E.8. We compare the daily values to each model’s forecasts using the Energy Score since these constitute multivariate data. We compute the score for all sliding windows of a given period, shown as boxplots. Aggregating the rainfall loses some spatial information, mostly retaining the specific characteristics of the region. On the other hand, using the Energy captures the temporal aspect.

Across all days and locations, the Cens-JGNM performs best, followed by the VAEGAN and ConvCNP. The spread of scores is lower for the Cens-JGNM while those of the two other models are larger.




Figure 7: **Cens-JGMM samples with fixed noise inputs:** We revert observed data to decorrelated Gaussian noise using an inverse Cholesky operation. We then use this noise to sample from our Cens-JGMM (bottom row) and compare these samples to observed data (top row) across 5 different days (each column). The likeness of samples to real data provides evidence of the good fit of our covariance matrix. It also highlights the stochasticity intrinsic to rainfall downscaling.

Further, for each model and time step, we compute the power spectrum of the rainfall field by applying a two-dimensional Fourier transform, squaring its magnitude, and azimuthally averaging over circular bins to obtain $P(k, t)$ at each wavenumber k . The ratio $P_{\text{pred}}(k, t)/P_{\text{obs}}(k, t)$ is then calculated for each time step and subsequently averaged over the entire 20-year prediction period to yield the final curves shown in Figure 6. In these plots, the x-axis represents the wavenumber in radians per meter (inversely related to the spatial scale), while the y-axis denotes the time-averaged power ratio, $\bar{R}(k)$. The horizontal line at $y = 1$ signifies an ideal match between predicted and observed spectra. Notably, the IFS (Hersbach et al., 2020), achieves the best agreement with the observed spectra across all wavenumbers. Among the neural network-based models, our proposed method Cens-JGMM demonstrates superior performance across the spectra, with an especially strong performance at high wavenumbers indicating a strong ability to effectively resolve fine-scale precipitation patterns—while both ConvCNP and VAEGAN show significant overestimation at lower wavenumbers (corresponding to larger spatial scales).

To test the correctness of our correlation matrix $\Sigma(\mathcal{D}|\theta)$, we generate samples from the Cens-JGMM using noise obtained from observed data. The idea is to use a Cholesky decomposition on the empirical covariance of observed data on the Gaussian scale to recover a decorrelated Gaussian version of the observed data. If the covariance is correctly specified, then the samples generated with the fixed observation-based noise as input will closely match the observed data. We detail this in Appendix E.10. We show the resulting samples in Figure 7 alongside observed data for five days in the test set.

For each of the days, the fixed noise samples from the Cens-JGMM resemble the observed data, with higher values of rain occurring in the same regions, and other patterns roughly being matched in each column. This suggests that the correlation matrix $\Sigma(\mathcal{D}|\theta)$ is well fitted to the data.

It more importantly exemplifies the inherent stochasticity of rainfall data as studied in this paper. As shown, in Figure 4, the samples obtained by our model are varied and do

not necessarily match the observed rain. However, marginal diagnostics suggest that our marginal models are well-calibrated, meaning they should correctly capture the distribution of observed rain. Now, with Figure 7, we have shown that the copula is also well estimated, since using the “correct” (obtained from observed data) standard Gaussian noise leads to samples that match observed data. Therefore, this suggests that observed rain is only a specific random realisation of standard noise, meaning the realised value could have been very different by pure chance; the observed data is only a glimpse of the randomness that could have been observed. This highlights the inherent stochasticity of this modelling problem, making probabilistic approaches that can capture this stochasticity essential.

We conclude our comparison with Table 3, presenting various diagnostics metrics assessing calibration, accuracy and spatial dependence. these metrics include the Continuous Ranked Probability Score (CRPS) [Appendix E.5], energy score [Appendix C], variogram score [Appendix E.6] as well as the Root Mean Square Error (RMSE) and Mean Absolute Bias (MAB) [Appendix E.7]. Our Cens-JGNM model outperforms the competing probabilistic downscaling models in all the diagnostics measures, with the exception of the CRPS where the ConvCNP is slightly better. The ConvCNP focuses all modelling efforts on its marginal models which we believe is the reason for its strong performance on the CRPS. For the RMSE and MAB, our Cens-JGNM stands a close second to the IFS benchmark. This is expected, as the IFS is a challenging benchmark to outperform here as this prediction is of higher resolution than the input variables used in our study.

Diagnostic Metrics for Probabilistic Forecasting				
Model	Cens-JGNM	ConvCNP	VAE-GAN	IFS
CRPS	3.5631	3.2678	4.1146	–
Energy Score	15.7253	17.0146	18.5991	–
Variogram Score	6,862,187	10,333,690	9,091,811	–
RMSE	3.8194	3.9314	4.4903	3.756
MAB	2.1438	2.2179	2.5820	2.118

Table 3: **Comparison with benchmarks:** Numerical diagnostics for our Cens-JGNM model against the ConvCNP, VAE-GAN models and the deterministic IFS. As IFS is a deterministic prediction, we can not compute probabilistic metrics for it. The Cens-JGNM generally outperforms other probabilistic models. For the RMSE and MAB, our model stands a close second to the IFS benchmark, which is a difficult benchmark to outperform due to its higher resolution than the inputs in our training data.

5.3 Robustness on Length of Training Data Set

In addition to comparing our approach against benchmarks, we appraise the robustness of the Cens-JGNM to reduced amounts of data. We train a Cens-JGNM on increasingly smaller data samples and compare the resulting forecasts on common test data. There are four Cens-JGNM versions in total, respectively fitted on 4 years of data (01/1985 to




Figure 8: **Robustness:** We show the (a) AUCs, (b) rank histograms, and (c) ESFs for the Cens-JGNM model versions fitted on decreasing amounts of data (10, 8, 6, and 4 years). The diagnostics show no trend across data size, indicating that the Cens-JGNM is robust to reduced amounts of data.

06/1989), 6 years (1983 to 06/1989), 8 years (1981 to 06/1989) and finally 10 years (1979 to 06/1989). All are tested on the same 30 years of data, from 07/1989 to 07/2019. We use the same diagnostics as in Sections 5.1 and 5.2 above, to determine whether the performance of our method suffers in settings with a lack of data.

We begin by inspecting the AUCs (with ROCs reported in Appendix F) in Figure 8a. The different sample sizes have no significant effect, as all four versions achieve similar AUCs. The models are better calibrated for lower rainfall events, which is explained by the reduced fitting data, implying that fewer extreme observations are used to train each model. We also compare their rank histograms in Figure 8b, which again show no significant difference or trend across model versions. From Figure 8c, showing ESFs, it is again indicated that periods above 4 years of data are sufficient to ensure the strong performance of our model as all four versions display nearly identical performance. Importantly, the performance of these versions matches the performance of the Cens-JGNM trained on 20 years of data in Section 5.1. Finally, in Table 4, we compare the metrics of the four model versions.

Diagnostic Metrics for Robustness Study				
Metric	4 years	6 years	8 years	10 years
CRPS	3.6129	3.4322	3.6996	3.5638
Energy Score	16.0627	16.0571	16.0603	15.9226
Variogram Score	4,918,176	4,690,858	5,069,616	5,102,977
RMSE	4.0362	4.0237	4.0275	3.9724
MAB	2.2336	2.2336	2.2295	2.1968

Table 4: **Robustness:** Numerical diagnostics for the Cens-JGNM approach fitted on multiple reduced amounts of training data. The results indicate no trend in the performance of our approach across training data quantity.

6. Related Work

In this section, we review some existing literature connected to our work. We begin in Section 6.1 by discussing other methods of forecasting rainfall in a univariate context. Next, in Section 6.2, we explore the estimation issues present in copulas when the likelihood becomes complicated to evaluate or directly unavailable. Finally, we describe alternative approaches to probabilistic rainfall forecasting than our two-step procedure in Section 6.3.

6.1 Marginal rainfall estimation

In Little et al. (2009), a study on GLMs used for rainfall density estimation with Markovian assumptions on temporal dependence recommends the use of a Bernoulli-Gamma mixture or an ensemble model, motivating our JGNM parametric choice. This comparison study further reinforces the findings of (Das, 1955; Fealy and Sweeney, 2007; Ailliot et al., 2015; Holsclaw et al., 2017; Bertolacci et al., 2019; Xie et al., 2023), all reporting success with marginal models based on Gamma densities. PCA regression was used in Li and Smith (2009) for precipitation downscaling using CM variables together with mean sea level pressure (MSLP) measurements, alleviating the underprediction problems faced by CMs. Our JGNM can be seen as a generalisation, where the refinement of variables used in a regression model is made in a data-driven way. Quantile regression is used by Tareghian and Rasmussen (2013) to get daily precipitation distributions by conditioning on CM outputs, increasing the flexibility in the selection of predictors, which our JGNM increases even further. Su et al. (2019) applied Bayesian model averaging (BMA) and stepwise regression models (SRMs) to downscaling. Our approach incorporates a pre-treatment of predictors as in SRMs with uncertainty quantification as in BMA, combining both strengths under a single approach.

Deep Learning (DL) methods have been used effectively in Geosciences, such as Li et al. (2021) applying DL methods to forecast location-wise densities of solar energy. In Adewoyin et al. (2021), expanding on previous DL models applied to rainfall forecasting such as Shi et al. (2015); Vandal et al. (2017); Miao et al. (2019) and Pan et al. (2019), the authors achieved state-of-the-art performance in point-wise prediction, showing that DL methods can be used to downscale rainfall based on CM forecasts. We directly improve on their approach by incorporating uncertainty quantification in the marginals.

6.2 Copula estimation problems

As is common when the copula is latent or in the presence of zero-inflated data (both being the case in our work), MLE-based inference becomes complicated or unavailable and alternatives need to be considered.

For extreme rainfall modelling, Huser and Wadsworth (2019) assume a censored copula for extreme values, targeting the full likelihood for inference, but only for up to 15 locations, noting that a high dimension is a limiting factor for Gaussian copulas, reinforcing the conclusions of Huser et al. (2017). Modelling extreme rainfall through a Brown–Resnick process, Richards and Wadsworth (2021) resort to a composite likelihood approach as an alternative to a computationally unfeasible MLE. More recently, in Richards et al. (2022) and Richards et al. (2023), authors use a censored Gaussian copula model for the dependence

of extreme rainfall events, circumventing the unavailable likelihood by adopting a pseudo-likelihood approach with spatially-informed sub-sampling. Outside of rainfall forecasting, DOBRA and LENKOSKI (2011) model graphical binary and ordinal variables in a Bayesian framework with a latent Gaussian copula, requiring an approximation to the likelihood function with the extended rank likelihood of Hoff (2007). Thus, as was summarised in Huser et al. (2016), while the use of these pseudo-likelihoods allows for inference, performing inference in higher dimensions remains an open problem. Our minimum scoring estimate approach provides a direct solution by not relying on likelihoods, additionally ensuring exact and computationally efficient parameter inference.

In the presence of zero-inflated data, Yoon et al. (2020) introduce a truncated latent Gaussian copula as an analogue to censored copulas and develop an estimator for the latent correlation matrix expanding on a rank-based procedure from Fan et al. (2017). This approach is expanded by Chung et al. (2022) to a Bayesian setting on graphical models for zero-inflated count-data data where inference is performed using Gibbs sampling. While their approach allows for higher dimensions, it is framed inside the Bayesian paradigm, whereas our minimum scoring rule approach is exact and can be used in both frequentist and Bayesian settings (with extensions following the works of Pacchiardi et al. (2024b)).

Recent advances in likelihood-free inference have led to the development of alternative methods for copulas inference. Alquier et al. (2022) use the maximum mean discrepancy (MMD) as a function to minimise in order to learn the parameters of copulas, obtaining a covariance estimator robust to miss-specification. Similarly, Janke et al. (2021) employ the energy distance (a possible instance of the MMD) as a loss function in the training of a generative network to approximate a copula distribution, avoiding the necessity for hyperparameters for the MMD. Our work extends their approaches to a spatial setting while accounting for censored observations.

6.3 Alternative probabilistic forecasting approaches for rainfall

Non-homogeneous hidden Markov models for rainfall occurrence are considered by Hughes et al. (1999) with extensions for amounts by Charles et al. (1999), modelling occurrence and amount of rainfall separately, which we do jointly with our JGNM. We further avoid their conditional independence assumptions on the spatial dependence structure with our censored copula approach. Frost et al. (2011), a review of six downscaling methods (namely Chandler (2002); Timbal (2004); McGregor (2005); Mehrotra and Sharma (2007); Chiew et al. (2009) and a HMMs-based (Charles et al., 1999) implementation of Kirshner (2007)), found that HMMs displayed weaknesses in capturing spatial correlation. Applying Bayesian HMMs to rainfall, Song et al. (2014) (expanding on the approach of Li and Smith (2009)) uses a covariance matrix construction with spatial distances used only in error correlation. Holsclaw et al. (2017) assume a more complicated time-dependency in transition probabilities, which calls for a Polya-Gamma-based MCMC sampling scheme. Bertolacci et al. (2019) use a spatial structure captured through a Gaussian process with marginal densities modelled by mixtures of expert models. Our approach extends these methods with direct and richer spatial modelling through the censored latent Gaussian copula.

Ensemble Model Output Statistics (EMOS), introduced by Gneiting et al. (2005) and improved by the works of Dabernig et al. (2017); Möller and Groß (2016, 2020), models

rainfall using an appropriate density of which parameters are estimated with linear regression equations incorporating spatio-temporal effects. Rasp and Lerch (2018) use neural networks instead of linear regression equations, further improved by Mlakar et al. (2024), performing estimation jointly across all future times and locations. Our approach similarly estimates distributions jointly in a spatio-temporal manner, however, we additionally employ the censored latent copula to impose the correct spatial dependence for any given forecast.

Machine learning approaches include Watson et al. (2020), with the use of Generative Adversarial Networks (GANs) to increase the resolution of rainfall forecasts while simultaneously aiming to correct biases. This work was expanded upon by Price and Rasp (2022) by using a conditional GAN based on coarse weather variables. Similarly, Harris et al. (2022) substitutes the GAN generator with a variational auto-encoder for conditional forecasts, adapted to cyclone-induced rainfall by Vosper et al. (2023). Alternatively, Pacchiardi et al. (2024a); Chen et al. (2024) use generative neural networks with the energy score as a loss function to avoid the problematic training of GANs. All these approaches learn an implicit likelihood and suffer from drawbacks such as the inability to generate zero values. Our JGNM model ensures explainability and a proper fitting of zero values and we compared our spatial dependence to the study of Harris et al. (2022), demonstrating improved performance. Further, our model can be evaluated on new locations after training, unlike most of these methods. Vaughan et al. (2022) employ a convolutional conditional neural process model with parametric marginals, explicit modelling of zero values and inclusion of out-of-sample locations. Their approach is very similar to the JGNM, although their model fails at capturing spatial dependence to the same extent as our censored copula method, as evidenced in Section 5.1. Ascenso et al. (2023) use a convolutional network approach with a modified loss function to enforce the spatial coherence of rainfall. We have a similar focus on spatial dependence with the energy score but additionally provide uncertainty quantification and interpretability.

7. Conclusion

In this paper, we introduced a novel method for spatio-temporal downscaling of zero-inflated data. We described our first component of the model, JGNM as a distributional learner capturing temporal trends in conjunction with surrounding information to describe distributions for quantities of interest at a location. The JGNM is very flexible and can be evaluated at any location, even one outside of the learning set. We also introduced censored latent Gaussian copulas as an answer to the spatial dependence modelling of censored data. We explained the intractability of the likelihood of this model, prohibiting the use of regular MLE-based methods. Hence we offered a scoring-rule-based inferential scheme, resulting in cheaper and more robust parameter inference. Our censored latent Gaussian copula methodology allows for the automatic inclusion of new locations, making our complete approach capable of accommodating additional locations added to the model without requiring parameter re-estimation. We compared our model to competitive benchmarks focusing on the same task (Vaughan et al., 2022; Harris et al., 2022). We demonstrated our model’s ability to correctly capture uncertainty to the same or better extent as competitors. Furthermore, we showcased our model’s superior capture of spatial dependence over

benchmarks, rendering our approach better suited for scenarios where spatial coherence is desired.

This work is an important improvement towards more reliable post-processing techniques as well as precipitation downscaling capable of correctly capturing uncertainty and spatial dependence. Moreover, while the main motivation of our work is rainfall downscaling, our methodology can effortlessly be applied to any spatio-temporal probabilistic prediction, with or without censoring. Therefore, we believe that our work is not only of interest to the weather science community but indeed relevant to the broader multivariate forecasting community. Further, we have used model fields as predictor variables, which can be resolved reasonably well by typical climate models, hence one possible future direction could be to use the developed methodology for post-processing of climate models.

Another possible extension of our model is the construction of the correlation matrix of the latent censored Gaussian copula in Section 4.1. Specifically, in the initial distance \mathcal{D} , one could include information on differences in weather variables to better capture the spatial dependence. In fact, it is possible to extend our copula approach to explicitly capture temporal dependence by adapting the kernel parameter according to temporal information. Other improvements particular to the modelling of extreme rainfall events could be made to the JGNM by considering extreme value distributions (Behrens et al., 2004; MacDonald et al., 2011; Ding et al., 2019; Gao et al., 2021; Richards et al., 2022), either entirely or by adding them as a mixture term in the tail. Finally, we acknowledge that imposing a Gaussian correlation structure, while computationally appealing and leading to easy spatial modelling, might be too restrictive, especially for extreme variations in rainfall levels at nearby locations. In this context, we could consider other types of copulas, no longer restricted to Gaussian dependence structures (Janke et al., 2021; Wadsworth and Tawn, 2022; Richards et al., 2023).

Acknowledgements and Disclosure of Funding

D.H. is funded by the Center for Doctoral Training in Mathematical Sciences at Warwick. R.A is funded by the University of Warwick and Southern University of Science and Technology. R.D. is funded by EPSRC (grant nos. EP/V025899/1 and EP/T017112/1) and NERC (grant no. NE/T00973X/1). The authors thank the reviewers for their helpful comments and the action editor for handling the revised submission.

Appendix A. Network Architecture

Our joint generalized neural model (JGNM) is designed to predict rainfall by parameterizing a three-parameter distribution. The model takes as input 6-hourly weather fields from the ERA5 reanalysis and predicts three parameters (p , μ and ϕ) for a predictive distributions for daily precipitation at a given point on a high-resolution grid.

Specifically, the input data consists of low-resolution weather fields

$$\mathbf{x}_j \in \mathbb{R}^{20 \times 21 \times 6}, \quad j = 1, \dots, 112,$$

where each \mathbf{x}_j represents one 6-hour time step over a 28-day period (4 time steps per day). These fields span a 20×21 grid covering the region of interest and include 6 weather variables

(specific humidity, air temperature, geopotential height at 500 hPa, eastward and northward wind components at 850 hPa, and total column water vapour). Note that precipitation is deliberately excluded from the input.

Following the approach in Adewoyin et al. (2021), we first perform bilinear interpolation to map the input from the coarse 20×21 grid to a finer 100×140 grid, matching the 0.1° resolution (approximately 8.5 km) of the observed rainfall data. Due to computational limits, we then partition the interpolated grid into 16×16 sub-sections which form the input to our neural network. For each sub-section, the 6-hourly data for all 6 variables is organized into a sequence of matrices

$$\mathbf{x}_j \in \mathbb{R}^{16 \times 16 \times 6}, \quad j = 1, \dots, 112.$$

Thus, the input for a 28-day period is represented as a sequence of shape (112, 16, 16, 6).

The model first applies a Time Distributed 2D Convolutional Layer (TD2L) that processes each \mathbf{x}_j independently with the same 2D convolution, transforming it into a hidden representation

$$\mathbf{h}_j^{\text{TD2L}} \in \mathbb{R}^{16 \times 16 \times 64}.$$

Next, the sequence of hidden representations is fed into a stacked Convolutional LSTM (ConvLSTM) module. This module consists of two bi-directional ConvLSTM layers with skip connections, which capture both spatial and temporal dependencies in the data. In our experiments, the ConvLSTM layers are configured with an input dropout of 0.25 and a recurrent dropout of 0.35.

Following the ConvLSTM, a Time Distributed Temporal Downscaling (TDTD) layer reduces the sequence length from 112 (6-hourly intervals) to 28 (daily values) by applying the same downscaling operation on non-overlapping groups of 4 time steps. The TDTD layer integrates a convolutional self-attention (CSA) module, a Gaussian Error Linear Unit (GeLU) activation (Hendrycks and Gimpel, 2016), and two 2D convolutional layers to effectively aggregate temporal information.

Finally, the output from the TDTD is passed through a Time-Distributed Mean Function Layer. This layer comprises three single-layer perceptrons, each acting as an inverse link function mapping the extracted features to one of the three distribution parameters. Consequently, the model generates parameter predictions:

$$\mu_i, \phi_i, p_i \in \mathbb{R}^{16 \times 16}, \quad i = 1, \dots, 28,$$

which represent daily predictive distributions of total rainfall at each location within the 16×16 grid, corresponding to the initial inputs $\mathbf{x}_j, j = 1, \dots, 112$. Finally, these spatially segmented 16×16 sub-images are combined to form a complete spatial prediction over the entire target region.

A schematic of the overall architecture is shown in Figure 9.

A.1 Convolutional Self-Attention (CSA)

The Convolutional Self-Attention (CSA) module is designed to reduce the temporal resolution of a sequence of three-dimensional feature maps, preserving maximal information




Figure 9: **Joint Generalised Neural Model (JGNM) architecture.** The JGNM takes as input a sequence of 112 tensors of shape $16 \times 16 \times 6$, corresponding to 6-hourly weather fields over 28 days. After initial processing by a Time-Distributed 2D Convolutional Layer (TD2L), the data passes through stacked bidirectional ConvLSTM layers capturing spatiotemporal dependencies. A Time-Distributed Temporal Downscaling Layer (TDTD) then aggregates temporal information from 6-hourly intervals to daily intervals, reducing sequence length from 112 to 28. Finally, a Time-Distributed Mean Function Layer outputs daily predictive parameters $p_i, \mu_i, \phi_i \in \mathbb{R}^{16 \times 16}$, parameterizing a predictive zero-gamma mixture distribution for total rainfall at each grid point. These 16×16 predictions are subsequently combined into the full spatial grid of size 100×140 .

from the original input while simultaneously incorporating long-range dependencies. In our implementation, the CSA module transforms an input tensor

$$x \in \mathbb{R}^{(b, t_{\text{input}}, h, w, c)},$$

where b is the batch size, t_{input} is the initial sequence length, h and w denote the spatial dimensions, and c is the number of channels, into an output tensor with a reduced temporal dimension $t_{\text{output}} = t_{\text{input}}/s$ (with s the downscaling factor).

Positional Embedding and Sub-sequence Formation: The input tensor is first partitioned into non-overlapping sub-sequences of length s . For each sub-sequence

$$x_{[j:j+s]} \in \mathbb{R}^{(b, s, h, w, c)},$$

we flatten the spatial dimensions to obtain

$$\tilde{x}_{[j:j+s]} \in \mathbb{R}^{(b, s, d)} \quad \text{with } d = h \times w \times c.$$

A learnable temporal positional embedding $P \in \mathbb{R}^{(1,s,d)}$ is then added to each sub-sequence to incorporate information about the order of the elements:

$$x'_{[j:j+s]} = \tilde{x}_{[j:j+s]} + P.$$

Query, Key, and Value Generation: For each sub-sequence $x'_{[j:j+s]}$, we compute the attention components as follows. The query vector is obtained by averaging the sub-sequence along the temporal dimension:

$$q_{[j:j+s]} = \frac{1}{s} \sum_{t=j}^{j+s-1} x'_t \in \mathbb{R}^{(b,d)}.$$

The key and value tensors are taken as the reshaped (but unaveraged) sub-sequence:

$$k_{[j:j+s]} = \tilde{x}_{[j:j+s]} \in \mathbb{R}^{(b,s,d)}, \quad v_{[j:j+s]} = \tilde{x}_{[j:j+s]} \in \mathbb{R}^{(b,s,d)}.$$

Multi-Head Projection: To enable the model to capture multiple types of relationships, we perform multi-head attention. For each head $h \in \{1, \dots, H\}$, the query, key, and value are linearly projected:

$$Q^h = q_{[j:j+s]} W_h^Q, \quad K^h = k_{[j:j+s]} W_h^K, \quad V^h = v_{[j:j+s]} W_h^V,$$

where $W_h^Q, W_h^K, W_h^V \in \mathbb{R}^{d \times d_{\text{hidden}}}$ are learned projection matrices. This results in:

$$Q^h \in \mathbb{R}^{(b,d_{\text{hidden}})}, \quad K^h, V^h \in \mathbb{R}^{(b,s,d_{\text{hidden}})}.$$

Scaled Dot-Product Attention: For each head, attention weights are computed using scaled dot-product attention:

$$\text{Attention}^h = \text{softmax} \left(\frac{Q^h (K^h)^\top}{\sqrt{d_{\text{hidden}}}} \right) V^h,$$

where the softmax is applied over the s (temporal) dimension, yielding an output of shape $\mathbb{R}^{(b,d_{\text{hidden}})}$ for each head.

Aggregation and Output Projection: The outputs from all H heads are concatenated:

$$\text{Concat}(\text{head}_1, \dots, \text{head}_H) \in \mathbb{R}^{(b,H \cdot d_{\text{hidden}})},$$

and then projected with a learned matrix $W^O \in \mathbb{R}^{(H \cdot d_{\text{hidden}}, d_{\text{output}})}$ to produce the final sub-sequence representation:

$$x_{[j:j+s]}^{(\text{output})} = \text{Concat}(\text{head}_1, \dots, \text{head}_H) W^O.$$

Reshaping to Recover Spatial Dimensions: Finally, the output for each sub-sequence is reshaped back to the original spatial dimensions, resulting in a tensor

$$x^{(\text{output})} \in \mathbb{R}^{(b,t_{\text{output}},h,w,c)},$$

which represents the temporally downsampled sequence of feature maps.

In our implementation, we set $h = 4$, $w = 4$, $c = 6$, $s = 4$, $d_{\text{hidden}} = 64$, $d_{\text{output}} = 64$, and use $H = 4$ attention heads. As t_{input}

A.2 Time Distributed Mean Function Layer

[p13]

In Table 5, we list the inverse link functions (mean functions) used in the Time Distributed Mean Function Layer for the Zero-Gamma mixture density. Note that these functions serve as activation functions mapping the real-valued output of the neural network back to the target parameter space. For example, with reference to our link functions introduced in Section 3.1, the inverse of g_μ for μ is implemented as

$$g_\mu^{-1}(x) = (6.0 \cdot \text{ReLU}(x + 0.4)) + 1 \times 10^{-6},$$

which ensures $\mu \in [0, \infty)$ while maintaining numerical stability. A similar formulation holds for ϕ via the inverse of g_ϕ .

	Parameter Range	g^{-1} : Mean (Inverse Link) Function
p	$(0, 1)$	$\text{sigmoid}(x)$
μ	$[0, \infty)$	$(6.0 \cdot \text{ReLU}(x + 0.4)) + 1 \times 10^{-6}$
ϕ	$[0, \infty)$	$(6.0 \cdot \text{ReLU}(x + 0.5)) + 1 \times 10^{-6}$

Table 5: **GLM Mean Functions:** Inverse link (activation) functions used in the Time Distributed Mean Function Layer of the Zero-Inflated JGNM.

Appendix B. Construction of Distances Matrix

Within a latent Gaussian copula, one needs to estimate a correlation matrix Σ , which we construct as a function of an initial distance matrix \mathcal{D} . As we want the copula to be spatially informed, we consider the latitude and longitude of locations as well as the topography given by geopotential height at a given location. For a given parameter value $a \in \mathbb{R}_{\geq 0}$, we construct the distance matrix \mathcal{D}^a by computing for all pairs of locations i, j :

$$\mathcal{D}_{(i,j)}^a = \sqrt{|lat_i - lat_j|^2 + |lon_i - lon_j|^2 + |a \cdot (topo_i - topo_j)|^2}.$$

Here, a serves to adequately weigh the importance of topography on the resulting correlation matrix $\Sigma = k(\mathcal{D}^a|\theta)$. With $a \rightarrow 0$, our copula model relies on geographical distances only while $a \rightarrow \infty$ makes the dependence rely uniquely on topography.

In a preliminary experiment, we experimented with optimising a, θ jointly, as well as setting a at pre-defined values (chosen at $a = 0, 0.005$, and 0.03) to only optimise θ . We show the Energy score values for different setups in Table 6 computed using forecasts after each optimisation set-up.

Set-up	$a = 0.0$	$a = 0.005$	$a = 0.03$	$a = 0.07226$ (joint optimisation)
Energy Score	15.7253	15.7394	15.7382	15.7852

Table 6: Energy score values for different values of a




Figure 10: **Samples with varying levels of topographical dependence:** We optimise θ for various fixed values of a as well as a joint optimisation set-up, and show samples for two days row-wise. The samples with higher topographical dependence are no longer realistic and over-rely on topography.

The best Energy score was obtained from the set-up where the topography was ignored by setting $a = 0$. The joint optimisation was unstable and obtained the worst score among all setups. We believe the reason for this issue was the presence of collinearity between topography and geographical distances, as locations close by will also share similar elevations. We show samples for two days (each row) from each optimisation outcome (each column) and target rain for the day, below in Figure 10. The samples were obtained with fixed noise coming from a Cholesky decomposition on observed rain, see Appendix E.10.

The realism flaws of relying excessively on topography become obvious. For higher values of a , the regions with rainfall are no longer determined by geographic closeness but rather by sharing the same elevations. This is particularly identifiable for the areas of Scotland which as a mountainous region ends up having the model misrepresent the dependence of samples.

Appendix C. Scoring Rules as a Divergence

Assume we have observed data from distribution \mathcal{P}^* and want to choose a parameter θ which parameterises a second distribution \mathcal{P}^θ such that the two distributions are as close to each other as possible. A divergence D is then defined as a function of two distributions such that (i) $D(\mathcal{P}^* || \mathcal{P}^\theta) \geq 0$ and (ii) $D(\mathcal{P}^* || \mathcal{P}^\theta) = 0 \iff \mathcal{P}^* = \mathcal{P}^\theta$. Therefore, a divergence can be used to optimise a parameter to recover the best possible model \mathcal{P}^θ for the data generating distribution \mathcal{P}^* .

One possible choice of divergences is scoring rules (SRs). As defined in Gneiting and Raftery (2007), a scoring rule $S(\mathcal{P}, \mathbf{z})$ is a function between a distribution \mathcal{P}^θ and observed data \mathbf{z} as a realisation of a random variable $\mathbf{Z} \sim \mathcal{P}^*$. Then the *expected scoring rule* is defined as $S(\mathcal{P}, \mathbf{z}) := \mathbb{E}_{\mathbf{Z} \sim \mathcal{P}^*} S(\mathcal{P}^\theta, \mathbf{Z})$. The SR is termed *proper* if relative to a set of distributions \mathbf{P} , if the expected SR is minimised when $\mathcal{P}^* = \mathcal{P}^\theta$:

$$S(\mathcal{P}^*, \mathcal{P}^*) \leq S(\mathcal{P}^\theta, \mathcal{P}^*) \quad \forall \mathcal{P}^\theta, \mathcal{P}^* \in \mathbf{P}.$$

Furthermore, a SR is termed *strictly proper*, if the minimisation above is unique:

$$S(\mathcal{P}^*, \mathcal{P}^*) < S(\mathcal{P}^\theta, \mathcal{P}^*) \quad \forall \mathcal{P}^\theta, \mathcal{P}^* \in \mathbf{P} \text{ s.t. } \mathcal{P}^* \neq \mathcal{P}^\theta.$$

By considering the quantity $D_{SR}(\mathcal{P}^* || \mathcal{P}^\theta) := S(\mathcal{P}^*, \mathcal{P}^\theta) - S(\mathcal{P}^*, \mathcal{P}^*)$ for a strictly proper SR, one can see that it defines a divergence. Indeed, (i) is verified by the SR being proper and (ii) is verified by the additional requirement of being strictly proper. This permits the use of strictly proper SRs as divergences to perform inference on parameters of a distribution. For the choice of SR, we introduce the Energy Score as:

$$S_E(\mathcal{P}^\theta, \mathbf{z}) = 2 \cdot \mathbb{E}_{\mathbf{z}' \sim \mathcal{P}^\theta} \|\mathbf{z}' - \mathbf{z}\|_2^\beta - \mathbb{E}_{\mathbf{z}'_1, \mathbf{z}'_2 \sim \mathcal{P}^\theta} \|\mathbf{z}'_1 - \mathbf{z}'_2\|_2^\beta$$

where $\beta \in (0, 2)$ is a hyperparameter regulating the severity of the divergence for incorrect distributions \mathcal{P}^θ . The energy SR is a strictly proper SR for the class of \mathbf{P} such that $\mathbb{E}_{\mathbf{z}' \sim \mathcal{P}} \|\mathbf{z}'\|^\beta < \infty$, see Gneiting and Raftery (2007) and is also known as a rescaling of the energy distance mentioned in section 4.2. It is possible to obtain unbiased estimates of $S_E(\mathcal{P}^\theta, \mathbf{z})$ by repeatedly sampling \mathbf{z}' m times from \mathcal{P}^θ , as:

$$\hat{S}_E(\{\mathbf{z}'_j : 1 \leq j \leq m\}, \mathbf{z}) = \frac{2}{m} \sum_{j=1}^m \|\mathbf{z}'_j - \mathbf{z}\|_2^\beta - \frac{1}{m(m-1)} \sum_{\substack{j,k=1 \\ k \neq j}}^m \|\mathbf{z}'_j - \mathbf{z}'_k\|_2^\beta$$

where \mathbf{z} are observations and \mathbf{z}' are samples. As such, we have a method for inferring parameters of a distribution by comparing the observations to simulated samples and minimising D_{SR} using the unbiased estimate of the SR. This is equivalent to choosing $\theta^* = \arg \min_{\theta} \hat{S}_E(\mathcal{P}^\theta, \mathbf{z})$ since the first part of D_{SR} is constant in θ .

Appendix D. Optimisation procedure

Here, we detail the training of θ from Section 4 as done in practice in experiments.

We choose a Matern kernel with $\nu = 0.5$, as for this value its expression simplifies Guttorp and Gneiting (2006) into:

$$\Sigma(\mathcal{D}; \theta)_{(i,j)} = \exp(-\mathcal{D}_{(i,j)}/\theta)$$

letting us implement a custom version of the kernel in Pytorch (Paszke et al., 2019) which allows for efficient gradient-based optimisation of our copula model.

To initialise the optimisation, we choose a θ value which minimises the Frobenius norm of $\Sigma(\mathcal{D}; \theta) - \hat{\Sigma}$ where $\hat{\Sigma}$ is the empirical covariance matrix of the observed rainfall transformed to the Gaussian scale $z_{i,s} = \Phi^{-1}(F_{i,s}(y_{i,s} | \mathcal{H}_s))$. This serves as an initial guess with a somewhat similar correlation matrix to that of the censored data and led to a more stable optimisation in initial experiments. In particular, for certain combinations of \mathcal{D} and θ , the resulting correlation matrix had many entries close to zero, which caused vanishing gradients in the Scoring Rule loss. Choosing an initialisation in this way prevented this issue in our experiments.

Throughout our whole optimisation, we take a fixed standard Gaussian sample with 103 realisations over all locations of size $(103, \mathcal{I})$. We then compute the correlation matrix with a given θ from our optimisation step and perform a Cholesky decomposition on the correlation matrix. We then multiply the decomposed matrix with the fixed noise sample to obtain a sample from an $|\mathcal{I}|$ -dimensional Gaussian with covariance matrix $\Sigma(\mathcal{D}; \theta)$. We then compute the Scoring Rule loss over minibatches of 103 days by comparing observed rainfall on the Gaussian scale to the censored samples obtained through a Cholesky decomposition of the matrix, following Section 4.3. We perform 500 gradient steps with the Adam optimiser (Kingma and Ba, 2014), with a learning rate of 0.2 and with gradient clipping.

The resulting optimisation converges for all experiments with a smooth optimisation surface based on the trace of parameter estimates and Scoring rules throughout iterations. The optimisation is not challenging as the parameter θ is one-dimensional and takes about 3 hours on a single Intel Xeon Platinum 8260 (Cascade Lake) CPU.

Appendix E. Description of Diagnostics

Here we describe the different diagnostics used.

E.1 L1 Norm of matrix differences

To compare two matrices, we use the L_1 norm for matrices. This is simply the sum of absolute entry-wise differences between the matrices. For matrices A, B , it is expressed as:

$$\sum_{i=1, j=1}^{|\mathcal{I}|} |A_{(i,j)} - B_{(i,j)}|$$

E.2 AUC of calibration task

We seek to verify the accuracy of our model’s daily estimated probabilities for a given amount of rainfall compared to the observed amount. To achieve this, we assess the sensitivity of the detection of heavy precipitation. We convert our forecasting task into a classification task by asking the following question: “Will there be precipitation exceeding a given amount?” The answer to this question will depend on a probability threshold, which our model has to surpass in order to emit a signal, or rather, give a positive answer to the question. This threshold τ is subjective and can be chosen as any value in $[0, 1]$.

For a given threshold, a test that correctly predicts heavy rain (as in rain exceeding a given amount) is known as a true positive. But if the test predicts heavy rain on a day it did not occur, this is known as a false positive. By treating precipitation at each day and point in space as separate independent events, a true positive rate (the proportion of points and days with heavy rain which was correctly detected) and the false positive rate (the proportion of points and days with no heavy rain with a positive signal) can be obtained.

In order to detect heavy rain, we compare the threshold τ with $1 - F_{i,t}(q)$, that is 1 minus the CDF for location i at time t evaluated at the heavy rain value $q \in \mathbb{R}_{\geq 0}$. The higher the value of τ , the more mass under the PDF on the right of q we need in order to be confident enough to issue a positive signal. On the contrary, assuming $\tau = 0$, we then classify all forecasts irrespective of location, time, or predictors as a positive.

By changing this threshold τ , we can construct a receiver operating characteristic (ROC)

curve plotting the true and false positive rates at each of these thresholds. Different levels of precipitation q can be tested, for example, 5 mm for light rain up to 25 mm for extreme events. The area under the ROC curve (AUC) can be used to assess how well the prediction for different levels of precipitation in the face of uncertainty was captured by the model. The closer the AUC gets to 1, the more confidence we can have in the heavy rainfall prediction capabilities of the examined model.

E.3 Rank histogram

Another way of assessing the calibration of a density is through rank histograms. This consists of computing the percentile value associated with an observation under a model, the percentiles being called the observation ranks. For the Cens-JGNM and ConvCNP models, as we have explicit models for marginal densities for every location $i \in \mathcal{I}$, we can compute the ranks of observed rain directly using the model’s CDF corresponding to location i and day t . In the case of 0-valued observations, we sample uniformly a value from the censored range $[0, 1 - p_{i,t}]$ following the explicit modelling of 0. For the VAEGAN model, we sample m times from the model to obtain an empirical distribution $\hat{F}_{i,t}$ for days t and location i . Then, we can estimate $F_{i,t}(y_{i,t})$ as $\hat{F}_{i,t}(y_{i,t}) = \sum_{k=1}^m \mathbb{1}\{y_{i,t} > y_{i,t}^{(k)}\}$, that is the percentile of the observed rainfall with respect to our model’s simulations $\{y_{i,t}^{(k)} : k \in \{1, \dots, m\}\}$. We note that to break ties in ranks of 0-valued rain, for the VAE-GAN, we add a meteorologically insignificant amount of rain to observations and simulations, following the original VAEGAN paper (Harris et al., 2022). By obtaining the rank of observations across days, we can construct a histogram of ranks. If the model is perfectly calibrated, the ranks will be uniformly distributed on $[0, 1]$, as should be the case for a CDF of its own realisations. We can assess our model’s properties by inspecting the dissimilarities of the histogram to a flat line (Hamill, 2001). For instance, if the model is over-dispersed, the ranks will tend to group in the middle while under-dispersion will manifest itself as most ranks falling within the extremes of the histogram. Equally, one can determine the bias of the model by observing unequal proportions of ranks on either side of the histogram, where an agglomeration on the right is indicative of under-prediction while agglomeration on the left indicates over-prediction.

E.4 Empirical Survival Function (ESF)

To further assess the calibration of our model, we test whether the estimated frequency of a type of event (eg. precipitation $> x$) over the testing period has matched the observed frequency. The estimated frequency of achieving a rainfall realisation above a level $x \in \mathbb{R}_{\geq 0}$ is obtained by generating samples from our model and looking at the proportion of samples exceeding x . This is then compared to the actual observation frequencies for values higher than x . A well-calibrated model’s sample frequency should follow the observation line closely, while a poorly calibrated model will deviate significantly from it. We refer to these plots as empirical survival function plots as they correspond to an approximation of $1 - F_{i,s}(y_{i,s} | \mathcal{H}_s)$.

E.5 CRPS

We also rely on numerical metrics to evaluate the effectiveness of our approach. Firstly, we consider the continuous ranked probability score (CRPS) (Matheson and Winkler, 1976), written as

$$\text{CRPS}(F_{i,t}, y_{i,t}) = - \int_{-\infty}^{\infty} (F(y'_{i,t}) - \mathbf{1}\{y'_{i,t} \geq y_{i,t}\})^2 dy'_{i,t}$$

which can be shown to be equal to

$$\text{CRPS}(F_{i,t}, y_{i,t}) = \frac{1}{2} E_{Y_1, Y_2 \sim F_{i,t}} |Y_1 - Y_2| - E_{Y \sim F_{i,t}} |Y - y_{i,t}|.$$

The above expectations can be estimated without bias through repeated sampling from the given model. The CRPS is a generalisation of the mean absolute error and indeed simplifies to it when considering a model with point-wise predictions. This metric is applied location-wise and favours well-calibrated forecasts.

E.6 Variogram score

To specifically target the spatial coherence of our model, we resort to the Variogram score, written as:

$$S_v^{(p)}(\mathbf{F}_t, \mathbf{y}_t) := \sum_{k,l=1}^n w_{kl} (|y_{k,t} - y_{l,t}|^p - \mathbb{E}_{\mathbf{Y} \sim \mathbf{F}_t} |Y_{k,t} - Y_{l,t}|^p)^2.$$

Again, we can appeal to repeated sampling of n samples in order to estimate the expectation. We take $p = 1$ for simplicity. The weights w_{kl} are chosen to represent the spatial nature of the problem. As such, we choose $w_{k,l}$ to be 1 over the distance from location k to l and setting $w_{kk} = 0$. This score takes into account the whole joint distribution and compares the smoothness of observed data against that of model forecasts. An ideal forecast would have a Variogram score of zero.

E.7 RMSE and MAB

Finally, we assess the models' performances by two quantitative metrics: the root mean squared error (RMSE) and mean absolute bias (MAB). In order to compute these quantities, we need to convert our probabilistic forecasts into point-forecasts. We do this by relying on the median of m forecasts for any given place and time. We can then compute the two metrics as:

$$\text{RMSE} = \sqrt{\frac{1}{n \times T} \sum_{i=1}^n \sum_{t=1}^T (y_{i,t} - \hat{y}_{i,t})^2}, \quad \text{MAB} = \frac{1}{n \times T} \sum_{i=1}^n \sum_{t=1}^T |y_{i,t} - \hat{y}_{i,t}|$$

where $\hat{y}_{i,t}$ is the median for time t and location i , $y_{i,t}$ is the recorded amount of rainfall at time t and location i , and T is the total number of days considered.

E.8 Spatio-Temporal Energy Score

To assess a model’s predictive performance for flood events, we test its effectiveness on aggregate rainfall prediction over a given region of locations $\mathcal{I}' \in \mathcal{I}$. Specifically, we compute the aggregate rain over all locations $i \in \mathcal{I}'$ for all days s in a period of days $s+ = (s, s+1, \dots)$, ranging from 1 to 7 days. This gives us one multivariate realization of aggregate daily rainfall for multiple days \mathbf{y}_{s+} , which we can compare against forecasted values using m model samples $\{y'_{s+,j}\}_{j=1}^m \sim \mathcal{P}_\theta$ with the Energy Score of Appendix C (in the case of 1 day, this reduces to the CRPS (Gneiting et al., 2005)). We compute this energy score for all sliding windows of $s+$ days in our experiments. The expression is as follows:

$$SR_{\text{spa-temp}}(\mathcal{P}_\theta, \mathbf{y}_{s+}) := \hat{S}_E(\{y'_{s+,j}\}_{j=1}^m, \mathbf{y}_{s+}).$$

E.9 Spectral Energy Plot

Mathematically, we define the time-averaged spectral ratio as

$$\bar{R}(k) = \frac{1}{T} \sum_{t=1}^T \left| \frac{P_{\text{pred}}(k, t)}{P_{\text{obs}}(k, t)} \right|,$$

where T is the total number of time steps. Here, the power spectrum $P(k, t)$ for a field $f(x, y, t)$ is computed from its 2D Fourier transform:

$$P(k, t) = \left\langle |\mathcal{F}\{f(x, y, t)\}|^2 \right\rangle_{\sqrt{k_x^2 + k_y^2 = k}},$$

with the averaging $\langle \cdot \rangle$ taken over all Fourier components corresponding to the wavenumber k . This concise formulation quantifies the average proportional difference in spectral energy between the predicted and observed fields across spatial scales.

E.10 Samples with fixed decorrelated noise.

Here we describe the experiment shown in Figure 7 of the main text. We begin by transforming observed data to the Gaussian scale, using

$$z_{i,s} = \Phi^{-1}(F_{i,s}(y_{i,s} | \mathcal{H}_s)).$$

We then compute the empirical covariance of these $\mathbf{z}_s = (z_{1,s}, \dots, z_{n,s})$ observations, call it $\hat{\Sigma}$. Then, we compute $\tilde{\mathbf{z}} = \mathbf{z} \cdot (\text{Chol}(\hat{\Sigma}))^{-1}$ where *Chol* represents a Cholesky decomposition, which (assuming $\hat{\Sigma}$ is the correct covariance) will give decorrelated standard Gaussian data. Next, we use this $\tilde{\mathbf{z}}$ as an input to a Cholesky decomposition-based-sampling with our $\Sigma(\mathcal{D}|\theta)$, to obtain $\tilde{\mathbf{u}} = \Phi(\tilde{\mathbf{z}} \cdot \text{Chol}(\Sigma(\mathcal{D}|\theta)))$, which (assuming $\tilde{\mathbf{z}}$ is standard Gaussian) will be a sample from the latent Gaussian copula with correlation matrix $\Sigma(\mathcal{D}|\theta)$. We then transform this back to the data scale with $\tilde{y}'_{i,s} = F_{i,s}^{-1}(u_{i,s})$.

Assuming that $\hat{\Sigma}$ is correct, this generated sample $\tilde{y}'_{i,s}$ should match closely the observed rainfall $y_{i,s}$ across all i . If $\Sigma(\mathcal{D}|\theta)$ and the covariance of the data closely match, it is straightforward to see that they should cancel in the transformations used above, meaning the generated samples should be identical. As such, the closer the samples, the more evidence there is that the correlation matrix of the copula model is correctly fit.

Appendix F. ROC for Experiments




Figure 11: **Benchmarking:** ROCs for two benchmark methods, namely the (a) the VAE-GAN and (b) ConvCNP as well as our (c) Cens-JGNM approach. Our model demonstrates the best performance, closely matched by the ConvCNP while the VAE-GAN achieves the lowest performance of the three methods.




Figure 12: **Robustness:** ROCs for robustness experiments where our Cens-JGNM approach was fitted on (a) 10, (b) 8, (c) 6 and (d) 4 years of data. No clear trend can be noticed, suggesting our approach is robust to any data amount size above four years.

Appendix G. Code Reproducibility

The code for the joint generalised neural model can be found at <https://github.com/Rilwan-Adewoyin/NeuralGLM>. Code for all figures shown and for the censored latent Gaussian copula can be found at https://github.com/Huk-David/SpaDep_VCop/tree/4Paper.

References

Rilwan A. Adewoyin, Peter Dueben, Peter Watson, Yulan He, and Ritabrata Dutta. TRUNET: a deep learning approach to high resolution prediction of rainfall. *Machine Learning*, 110(8):2035–2062, Jul 2021. ISSN 1573-0565. doi: 10.1007/s10994-021-06022-6. URL <http://dx.doi.org/10.1007/s10994-021-06022-6>.

- Pierre Ailliot, Denis Allard, Valérie Monbet, and Philippe Naveau. Stochastic weather generators: an overview of weather type models. *Journal de la société française de statistique*, 156(1):101–113, 2015.
- Pierre Alquier, Badr-Eddine Chérif-Abdellatif, Alexis Derumigny, and Jean-David Fermanian. Estimation of copulas via maximum mean discrepancy. *Journal of the American Statistical Association*, pages 1–16, 2022.
- Guido Ascenso, Andrea Ficchi, Leone Cavicchia, Enrico Scoccimarro, Matteo Giuliani, Andrea Castelletti, et al. Improving the spatial accuracy of extreme tropical cyclone rainfall in ERA5 using deep learning. In *EGU General Assembly Conference Abstracts*, pages EGU–8085, 2023.
- Cibele N Behrens, Hedibert F Lopes, and Dani Gamerman. Bayesian analysis of extreme events with threshold estimation. *Statistical modelling*, 4(3):227–244, 2004.
- Michael Bertolacci, Edward Cripps, Ori Rosen, John W Lau, and Sally Cripps. Climate inference on daily rainfall across the Australian continent, 1876–2015. *The Annals of Applied Statistics*, 13(2):683–712, 2019.
- RE Chandler. Glimclim: Generalized linear modelling for daily climate time series (software and user guide). 2002.
- Stephen P Charles, Bryson C Bates, and James P Hughes. A spatiotemporal model for downscaling precipitation occurrence and amounts. *Journal of Geophysical Research: Atmospheres*, 104(D24):31657–31669, 1999.
- Jieyu Chen, Tim Janke, Florian Steinke, and Sebastian Lerch. Generative machine learning methods for multivariate ensemble postprocessing. *The Annals of Applied Statistics*, 18(1):159–183, 2024.
- FHS Chiew, J Teng, J Vaze, DA Post, JM Perraud, DGC Kirono, and NR Viney. Estimating climate change impact on runoff across southeast Australia: Method, results, and implications of the modeling method. *Water Resources Research*, 45(10), 2009.
- Hee Cheol Chung, Irina Gaynanova, and Yang Ni. Phylogenetically informed Bayesian truncated copula graphical models for microbial association networks. *The Annals of Applied Statistics*, 16(4):2437–2457, 2022.
- R. Cornes, G. van der Schrier, E. van den Besselaar, and P. Jones. An ensemble version of the E-OBS temperature and precipitation datasets: version 21.0e, 2019.
- Markus Dabernig, Georg J Mayr, Jakob W Messner, and Achim Zeileis. Spatial ensemble post-processing with standardized anomalies. *Quarterly Journal of the Royal Meteorological Society*, 143(703):909–916, 2017.
- SC Das. The fitting of truncated type III curves to daily rainfall data. *Australian Journal of Physics*, 8(2):298–304, 1955.

- A Philip Dawid, Monica Musio, and Laura Ventura. Minimum scoring rule inference. *Scandinavian Journal of Statistics*, 43(1):123–138, 2016.
- Daizong Ding, Mi Zhang, Xudong Pan, Min Yang, and Xiangnan He. Modeling extreme events in time series prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1114–1122, 2019.
- Adrian DOBRA and Alex LENKOSKI. Copula gaussian graphical models and their application to modeling functional disability data. *The Annals of applied statistics*, 5(2): 969–993, 2011.
- James Duncan, Shashank Subramanian, and Peter Harrington. Generative modeling of high-resolution global precipitation forecasts. *arXiv preprint arXiv:2210.12504*, 2022.
- Peter K. Dunn and Gordon K. Smyth. *Generalized Linear Models With Examples in R*. Springer New York, 2018. doi: 10.1007/978-1-4419-0118-7. URL <https://doi.org/10.1007%2F978-1-4419-0118-7>.
- Jianqing Fan, Han Liu, Yang Ning, and Hui Zou. High dimensional semiparametric latent graphical model for mixed data. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 405–421, 2017.
- Rowan Fealy and John Sweeney. Statistical downscaling of precipitation for a selection of sites in Ireland employing a generalised linear modelling approach. *International Journal of Climatology: A Journal of the Royal Meteorological Society*, 27(15):2083–2094, 2007.
- Andrew J Frost, Stephen P Charles, Bertrand Timbal, Francis HS Chiew, Raj Mehrotra, Kim C Nguyen, Richard E Chandler, John L McGregor, Guobin Fu, Dewi GC Kirono, et al. A comparison of multi-site daily rainfall downscaling techniques under Australian conditions. *Journal of Hydrology*, 408(1-2):1–18, 2011.
- Jiajia Gao, Pengfei Ma, Jun Du, and Xiaoqing Huang. Spatial distribution of extreme precipitation in the Tibetan plateau and effects of external forcing factors based on generalized Pareto distribution. *Water Supply*, 21(3):1253–1262, 2021.
- Florian Gerber and Douglas Nychka. Fast covariance parameter estimation of spatial gaussian process models using neural networks. *Stat*, 10(1):e382, 2021.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437. URL <https://doi.org/10.1198/016214506000001437>.
- Tilmann Gneiting, Adrian E Raftery, Anton H Westveld, and Tom Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133(5):1098–1118, 2005.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.

- Peter Guttorp and Tilmann Gneiting. Studies in the history of probability and statistics xlix on the matérn correlation family. *Biometrika*, 93(4):989–995, 2006.
- Thomas M Hamill. Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129(3):550–560, 2001.
- Lucy Harris, Andrew T. T. McRae, Matthew Chantry, Peter D. Dueben, and Tim N. Palmer. A generative deep learning approach to stochastic downscaling of precipitation forecasts. *Journal of Advances in Modeling Earth Systems*, 14(10), oct 2022. doi: 10.1029/2022ms003120. URL <https://doi.org/10.1029/2022ms003120>.
- Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs), 2016.
- Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, and Joaquín Muñoz-Sabater. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: 10.1002/qj.3803.
- Peter D Hoff. Extending the rank likelihood for semiparametric copula estimation. *The Annals of Applied Statistics*, pages 265–283, 2007.
- Tracy Holsclaw, Arthur M Greene, Andrew W Robertson, and Padhraic Smyth. Bayesian nonhomogeneous markov models via pólya-gamma data augmentation with applications to rainfall modeling. *The Annals of Applied Statistics*, 11(1), 2017.
- James P Hughes, Peter Guttorp, and Stephen P Charles. A non-homogeneous hidden markov model for precipitation occurrence. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 48(1):15–30, 1999.
- Raphaël Huser and Jennifer L Wadsworth. Modeling spatial processes with unknown extremal dependence class. *Journal of the American Statistical Association*, 114(525):434–444, 2019.
- Raphaël Huser, Anthony C Davison, and Marc G Genton. Likelihood estimators for multivariate extremes. *Extremes*, 19:79–103, 2016.
- Raphaël Huser, Thomas Opitz, and Emeric Thibaud. Bridging asymptotic independence and dependence in spatial extremes using Gaussian scale mixtures. *Spatial Statistics*, 21: 166–186, 2017.
- Tim Janke, Mohamed Ghanmi, and Florian Steinke. Implicit generative copulas. *Advances in Neural Information Processing Systems*, 34:26028–26039, 2021.
- Thomas Jung, Gianpaolo Balsamo, Peter Bechtold, ACM Beljaars, Martin Koehler, MJ Miller, J-J Morcrette, Andrew Orr, MJ Rodwell, and Adrian M Tompkins. The ECMWF model climate: recent progress through improved physical parametrizations. *Quarterly Journal of the Royal Meteorological Society*, 136(650):1145–1160, 2010.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Sergey Kirshner. Learning with tree-averaged densities and distributions. *Advances in Neural Information Processing Systems*, 20, 2007.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- Simon Lang, Mihai Alexe, Matthew Chantry, Jesper Dramsch, Florian Pinault, Baudouin Raoult, Mariana CA Clare, Christian Lessig, Michael Maier-Gerber, Linus Magnusson, et al. Aifs–ecmwf’s data-driven forecasting system. *arXiv preprint arXiv:2406.01465*, 2024.
- David A. Lavers, Adrian Simmons, Freja Vamborg, and Mark J. Rodwell. An evaluation of era5 precipitation for climate monitoring. *Quarterly Journal of the Royal Meteorological Society*, 148(748):3152–3165, 2022. doi: <https://doi.org/10.1002/qj.4351>. URL <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.4351>.
- Rui Li, Brian J Reich, and Howard D Bondell. Deep distribution regression. *Computational Statistics & Data Analysis*, 159:107203, 2021.
- Yun Li and Ian Smith. A statistical downscaling model for Southern Australia winter rainfall. *Journal of Climate*, 22(5):1142 – 1158, 2009. doi: 10.1175/2008JCLI2160.1. URL <https://journals.ametsoc.org/view/journals/clim/22/5/2008jcli2160.1.xml>.
- Max A Little, Patrick E McSharry, and James W Taylor. Generalized linear models for site-specific density forecasting of uk daily rainfall. *Monthly Weather Review*, 137(3): 1029, 2009.
- Anna MacDonald, Carl John Scarrott, Dominic Lee, Brian Darlow, Marco Reale, and Glynn Russell. A flexible extreme value mixture model. *Computational Statistics & Data Analysis*, 55(6):2137–2157, 2011.
- Bertil Matérn. *Spatial variation*, volume 36. Springer Science & Business Media, 2013.
- James E Matheson and Robert L Winkler. Scoring rules for continuous probability distributions. *Management science*, 22(10):1087–1096, 1976.
- John L McGregor. *C-CAM: Geometric aspects and dynamical formulation*, volume 70. CSIRO Atmospheric Research Dickson ACT, Australia, 2005.
- R Mehrotra and Ashish Sharma. Preserving low-frequency variability in generated daily rainfall sequences. *Journal of Hydrology*, 345(1-2):102–120, 2007.
- Qinghua Miao, Baoxiang Pan, Hao Wang, Kuolin Hsu, and Soroosh Sorooshian. Improving monsoon precipitation prediction using combined convolutional and long short term memory neural network. *Water*, 11(5):977, 2019.
- Peter Mlakar, Janko Merše, and Jana Faganeli Pucer. Ensemble weather forecast post-processing with a flexible probabilistic neural network approach. *Quarterly Journal of the Royal Meteorological Society*, 150(764):4156–4177, 2024.

- Annette Möller and Jürgen Groß. Probabilistic temperature forecasting based on an ensemble autoregressive modification. *Quarterly Journal of the Royal Meteorological Society*, 142(696):1385–1394, 2016.
- Annette Möller and Jürgen Groß. Probabilistic temperature forecasting with a heteroscedastic autoregressive ensemble postprocessing model. *Quarterly Journal of the Royal Meteorological Society*, 146(726):211–224, 2020.
- John A Nelder and Youngjo Lee. Generalized linear models for the analysis of Taguchi-type experiments. *Applied stochastic models and data analysis*, 7(1):107–120, 1991.
- Roger B Nelsen. *An introduction to copulas*. Springer, 2006.
- Lorenzo Pacchiardi and Ritabrata Dutta. Score matched neural exponential families for likelihood-free inference. *Journal of Machine Learning Research*, 23(38):1–71, 2022a.
- Lorenzo Pacchiardi and Ritabrata Dutta. Likelihood-free inference with generative neural networks via scoring rule minimization. *arXiv preprint arXiv:2205.15784*, 2022b.
- Lorenzo Pacchiardi, Rilwan A Adewoyin, Peter Dueben, and Ritabrata Dutta. Probabilistic forecasting with generative networks via scoring rule minimization. *Journal of Machine Learning Research*, 25(45):1–64, 2024a.
- Lorenzo Pacchiardi, Sherman Khoo, and Ritabrata Dutta. Generalized bayesian likelihood-free inference. *Electronic Journal of Statistics*, 18(2):3628–3686, 2024b.
- Baoxiang Pan, Kuolin Hsu, Amir AghaKouchak, and Soroosh Sorooshian. Improving precipitation estimation using convolutional neural network. *Water Resources Research*, 55(3):2301–2321, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Ilan Price and Stephan Rasp. Increasing the accuracy and resolution of precipitation forecasts using deep generative models. In *International conference on artificial intelligence and statistics*, pages 10555–10571. PMLR, 2022.
- Wenchao Qi, Chao Ma, Hongshi Xu, Zifan Chen, Kai Zhao, and Hao Han. A review on applications of urban flood models in flood mitigation strategies. *Natural Hazards*, 108:31–62, 2021.
- David A Randall, Richard A Wood, Sandrine Bony, Robert Colman, Thierry Fiechefet, John Fyfe, Vladimir Kattsov, Andrew Pitman, Jagadish Shukla, Jayaraman Srinivasan, et al. Climate models and their evaluation. In *Climate change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC (FAR)*, pages 589–662. Cambridge University Press, 2007.

- Carl Edward Rasmussen, Christopher KI Williams, et al. *Gaussian processes for machine learning*, volume 1. Springer, 2006.
- Stephan Rasp and Sebastian Lerch. Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, 146(11):3885–3900, 2018.
- Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, et al. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878):672–677, 2021.
- Jordan Richards and Jennifer L Wadsworth. Spatial deformation for nonstationary extremal dependence. *Environmetrics*, 32(5):e2671, 2021.
- Jordan Richards, Jonathan A Tawn, and Simon Brown. Modelling extremes of spatial aggregates of precipitation using conditional methods. *The Annals of Applied Statistics*, 16(4):2693–2713, 2022.
- Jordan Richards, Jonathan A Tawn, and Simon Brown. Joint estimation of extreme spatially aggregated precipitation at different scales through mixture modelling. *Spatial Statistics*, page 100725, 2023.
- Jordan Richards, Matthew Sainsbury-Dale, Andrew Zammit-Mangion, and Raphaël Huser. Neural bayes estimators for censored inference with peaks-over-threshold models. *Journal of Machine Learning Research*, 25(390):1–49, 2024.
- Justin T Schoof. Statistical downscaling in climatology. *Geography Compass*, 7(4):249–265, 2013.
- Xingjian Shi, Zhourong Chen, and Hao Wang. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems 28*, pages 802–810. Curran Associates, Inc., 2015.
- Kunio Shimizu. A bivariate mixed lognormal distribution with an analysis of rainfall data. *Journal of Applied Meteorology and Climatology*, 32(2):161–171, 1993.
- M Sklar. Fonctions de répartition à n dimensions et leurs marges. In *Annales de l’ISUP*, volume 8, pages 229–231, 1959.
- Yong Song, Yun Li, Bryson Bates, and Christopher K Wikle. A Bayesian hierarchical downscaling model for south-west Western Australia rainfall. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 63(5):715–736, 2014.
- Haifeng Su, Zhe Xiong, Xiaodong Yan, and Xingang Dai. An evaluation of two statistical downscaling models for downscaling monthly precipitation in the Heihe River basin of China. *Theoretical and Applied Climatology*, 138:1913–1923, 2019.
- Reza Tareghian and Peter Rasmussen. Statistical downscaling of precipitation using quantile regression. *Journal of Hydrology*, 487:122–135, 04 2013. doi: 10.1016/j.jhydrol.2013.02.029.

- Bertrand Timbal. Southwest Australia past and future rainfall trends. *Climate Research*, 26(3):233–249, 2004.
- M.-N. Tran, N. Nguyen, D. Nott, and R. Kohn. Bayesian deep net GLM and GLMM. *Journal of Computational and Graphical Statistics*, 29(1):97–113, 2020. doi: 10.1080/10618600.2019.1637747. URL <https://doi.org/10.1080/10618600.2019.1637747>.
- Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew Michaelis, Ramakrishna Nemani, and Auroop R. Ganguly. DeepSD: Generating high resolution climate change projections through single image super-resolution. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 1663–1672, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450348874. doi: 10.1145/3097983.3098004.
- Thomas Vandal, Evan Kodra, Jennifer Dy, Sangram Ganguly, Ramakrishna Nemani, and Auroop R. Ganguly. Quantifying uncertainty in discrete-continuous and skewed data with Bayesian deep learning. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Jul 2018. doi: 10.1145/3219819.3219996.
- Anna Vaughan, Will Tebbutt, J Scott Hosking, and Richard E Turner. Convolutional conditional neural processes for local climate downscaling. *Geoscientific Model Development*, 15(1):251–268, 2022.
- Emily Vosper, Peter Watson, Lucy Harris, Andrew McRae, Raul Santos-Rodriguez, Laurence Aitchison, and Dann Mitchell. Deep learning for downscaling tropical cyclone rainfall to hazard-relevant spatial scales. *Journal of Geophysical Research: Atmospheres*, page e2022JD038163, 2023.
- Jennifer L Wadsworth and JA Tawn. Higher-dimensional spatial extremes via single-site conditioning. *Spatial Statistics*, 51:100677, 2022.
- Pascaline Wallemacq and Claudia Herden. *The Human Cost Of Weather Related Disasters*. CRED, UNIDSR, 2015.
- Fang Wang, Di Tian, and Mark Carroll. Customized deep learning for precipitation bias correction and downscaling. *Geoscientific Model Development*, 16(2):535–556, 2023.
- Campbell Watson, Chulin Wang, Tim Lynar, and Kommy Weldemariam. Investigating two super-resolution methods for downscaling precipitation: Esrgan and car. 2020.
- Peter AG Watson. Machine learning applications for weather and climate need greater focus on extremes. *Environmental Research Letters*, 17(11):111004, 2022.
- Wang-Chun Woo and Wai-Kin Wong. Operational application of optical flow techniques to radar-based rainfall nowcasting. *Atmosphere*, 8(3), 2017. ISSN 2073-4433. doi: 10.3390/atmos8030048.
- Li Xiang, Jie Xiang, Jiping Guan, Fuhan Zhang, Yanling Zhao, and Lifeng Zhang. A novel reference-based and gradient-guided deep learning model for daily precipitation downscaling. *Atmosphere*, 13(4):511, 2022.

Kun Xie, Yanfeng He, Jong-Suk Kim, Sun-Kwon Yoon, Jie Liu, Hua Chen, Jung Hwan Lee, Xiang Zhang, and Chong-Yu Xu. Assessment of the joint impact of rainfall characteristics on urban flooding and resilience using the copula method. *Water Resources Management*, 37(4):1765–1784, 2023.

Grace Yoon, Raymond J Carroll, and Irina Gaynanova. Sparse semiparametric canonical correlation analysis for data of mixed types. *Biometrika*, 107(3):609–625, 2020.