

Spatiotemporally Consistent Multivariate Bias Correction for Climate Projections via Nested Vine Copulas

Theresa Meier^{1,2,3,*}, Erwan Koch^{1,2}, Valérie Chavez-Demoulin^{1,2} and Thibault Vatter^{1,2,3}

March 13, 2026

¹ Faculty of Business and Economics (HEC), University of Lausanne, Switzerland

² Expertise Center for Climate Extremes, Faculty of Business and Economics - Faculty of Geosciences and Environment, University of Lausanne, Switzerland

³ University of Applied Sciences and Arts Western Switzerland (HES-SO), Geneva, Switzerland

Abstract

Climate models are essential for understanding large-scale climate dynamics and long-term climate change, yet they exhibit systematic biases when compared with historical observations. Existing multivariate bias correction (MBC) approaches do not explicitly handle spatiotemporal dependence. However, preserving both spatiotemporal and inter-variable consistency is essential for realistic climate dynamics and reliable regional impact assessments. To address this gap, we propose a novel MBC method called GN-VBC that uses generalized additive models (GAMs) to disentangle spatiotemporal deterministic effects from stochastic residuals. To model joint distributions and dependencies across variables and locations, we introduce nested vine copulas (NVCs), a hierarchical vine merging strategy. NVC in the context of MBC combines two dependence levels: (i) spatial dependence across locations, modeled separately for each variable, and (ii) inter-variable dependence modeled at a selected reference location, which links the spatial models into a coherent multivariate and spatial structure. An application to Switzerland shows improvements in preserving inter-variable, spatial and temporal dependence across a wide range of evaluation metrics.

Keywords: Bias correction; Climate projections; Generalized additive models; Nested vine copulas; Spatiotemporal consistency.

*The authors gratefully acknowledge Sven Kotlarski (MeteoSwiss) for providing Swiss data products, the CORDEX initiative, and ECMWF for providing ERA5-Land data.

1 Introduction

Global climate models (GCMs) provide physically consistent simulations of the Earth system and remain the foundation for understanding large-scale climate dynamics and long-term climate change (e.g., [Flato et al. 2013](#)). Their coarse spatial resolution, however, limits their ability to capture regional processes and extremes. Regional climate models (RCMs) complement GCMs by dynamically downscaling large-scale climate information, adding finer spatial detail and a more realistic representation of mesoscale processes such as topography-driven precipitation and land–atmosphere feedbacks (e.g., [Rummukainen 2016](#)). Together, GCMs and RCMs provide complementary climate information that supports the assessment of changes in climate hazards and extreme events. Although the resulting projections remain too coarse for local-scale impact studies, they constitute the essential basis for subsequent downscaling and impact-oriented modeling frameworks used in adaptation planning.

However, GCMs and RCMs exhibit biases, i.e., systematic differences between simulated climate statistics (such as a mean, quantile, or event frequency) and the corresponding statistics computed from observations or from a reference dataset (e.g., [Ehret et al. 2012](#)). They arise from limited spatial resolution, simplified physical processes, and incomplete knowledge of the climate system. Typically, GCMs and RCMs often exhibit errors in mean conditions, variability, seasonal cycles, extremes, and spatiotemporal dependence, which motivate the widespread use of bias correction before applying their outputs in impact assessments and risk analyses. For instance, temperature biases of up to $\pm 1.5^\circ\text{C}$ and precipitation biases reaching $\pm 40\%$ have been reported over Europe (e.g., [Kotlarski et al. 2014](#)), while multi-model ensembles often show systematic regional biases in precipitation, temperature, or wind (e.g., [Vautard et al. 2021](#)).

To formalize the notion of bias correction, let Y denote a climate variable. We use su-

perscripts m and r to refer to model and reference data, and c and p to denote the range of timestamps corresponding to the calibration and projection periods. For example, $Y^{(mc)} = \{Y_t^{(m)} : t \in c\}$ denotes model data during the calibration period. Reference observations are available only during the calibration period, whereas model simulations exist in both periods. In our application, however, we select a projection period for which reference data are also available, allowing us to evaluate our approach. A bias correction is a transformation h applied to model output such that the corrected values $\tilde{Y}_t^{(m)} = h(Y_t^{(m)})$ yield projected series $\tilde{Y}^{(mp)}$ whose distribution resembles that of the (unobserved) reference data during the projection period, $Y^{(rp)}$. In the univariate stationary setting, this reduces, e.g., to quantile mapping (QM; e.g., [Rajczak et al. 2016](#), [Ivanov & Kotlarski 2017](#)), where $h(\cdot) = \hat{Q}^{(rc)} \circ \hat{F}^{(mc)}(\cdot)$ with $\hat{Q}^{(rc)}$ and $\hat{F}^{(mc)}$ the empirical quantile and distribution functions of the reference and model data during the calibration period, respectively. In practice, however, since climate processes are inherently multivariate and spatially structured, bias correction must address more than univariate discrepancies.

The aim of this work is to introduce a general bias correction method that explicitly and simultaneously adjusts inter-variable, spatial, and temporal dependence. We apply the method to a challenging real-world setting in Switzerland ([Figure 1](#)),

where complex Alpine topography imposes strong spatial heterogeneity. The application involves jointly bias-correcting the five key atmospheric variables of [Table 1](#), projected within the Coordinated Regional Climate Downscaling Experiment for the European domain (EURO-CORDEX) at 22 grid points. [Figure 2](#) shows that the climate model (see [Section 5.1](#) for a

Abbrev.	Atmospheric Variable
tas	Mean 2-m temperature (K)
pr	Total precipitation (mm)
hurs	Mean 2-m relative humidity
sfcWind	Mean 10-m wind speed (m s^{-1})
ps	Mean surface pressure (Pa)

Table 1: Variables in the case study.

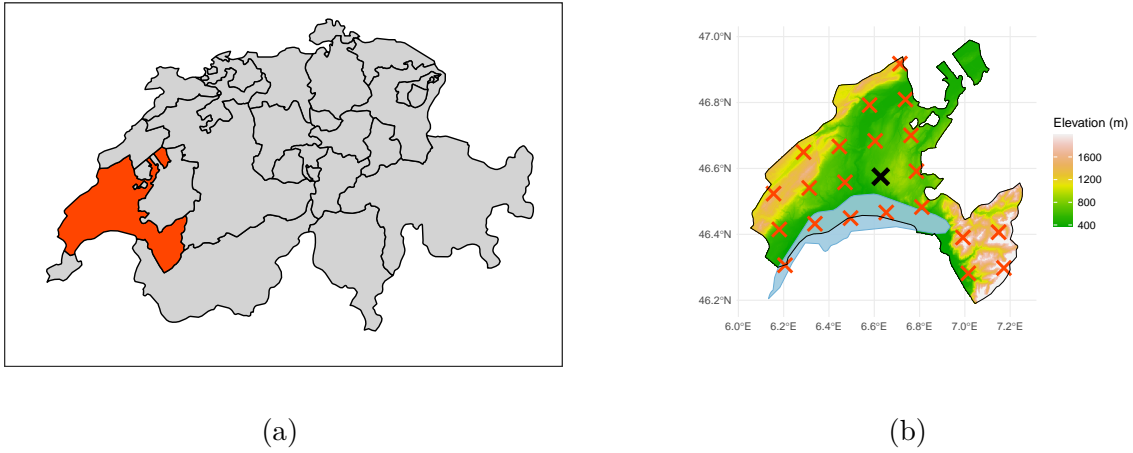


Figure 1: Panel (a) highlights the canton of Vaud in Switzerland. Panel (b) shows the 22 grid points in the study, with the black cell serving as the bridging location in [Section 5](#).

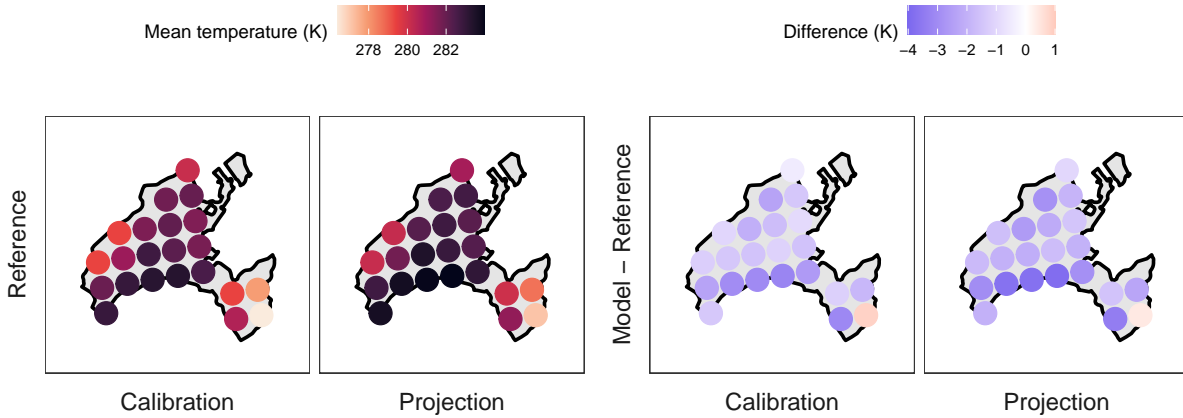


Figure 2: Mean reference temperature per grid point for the calibration (1980–2009) and projection (2010–2022) periods (left), and corresponding model-reference differences (right).

detailed description) underestimates mean temperatures across most locations. While this reflects a univariate bias (per variable and per location), further biases arise in inter-variable dependence structures and spatial coherence. Our objective in the application ([Section 5](#)) is therefore to obtain bias-corrected fields that faithfully reproduce reality in terms of multivariate and spatiotemporal dependence.

A wide range of bias correction methods has been proposed to enhance the fidelity of climate model outputs. Early univariate approaches relied on delta or scaling adjustments, which

correct mean or variance biases while preserving modeled change signals. More flexible distribution-based approaches were later developed, most notably quantile mapping and its extensions such as detrended quantile mapping (DQM) and quantile delta mapping (QDM), which aim to preserve long-term climate trends while correcting distributional biases (Cannon et al. 2015, Rajczak et al. 2016, Ivanov & Kotlarski 2017). These methods, however, operate on each variable and each location independently, and several multivariate bias correction (MBC) methods, that seek to correct biases in the joint distributions, have been proposed. State-of-the-art approaches include correlation-based methods (Bárdossy & Pegram 2012, Cannon 2016), the normalization–rotation algorithm MBCn (Cannon 2018), the rank-resampling method R2D2 (Vrac 2018), and optimal-transport-based corrections such as dOTC (Robin et al. 2019); see François et al. (2020) for a review. Nevertheless, existing methods remain limited in their ability to simultaneously reproduce inter-variable, spatial, and temporal dependence. In addition, the underlying dependence structures are often not represented explicitly, and the three types of dependence are typically handled in the same way, which reduces flexibility and interpretability.

Closer to our work, the vine copula bias correction (VBC) framework (Funk et al. 2025) uses vine copulas to model complex dependence structures and the Rosenblatt transform (Rosenblatt 1952) for a multivariate analogue of quantile mapping. Nonetheless, the method is applied location-wise and does not account for seasonality. Beyond bias correction, vine copulas have also been used to model spatial dependence in environmental data. For example, Gräler (2014) constructs distance-based vine models for spatial extremes, while Erhardt et al. (2015a) and Erhardt et al. (2015b) introduce vine structures parametrized by spatial covariates such as geographic distance and elevation. However, those approaches typically focus on a single variable across locations and therefore do not simultaneously address spatial and inter-variable dependence.

To address the limitations of existing methods, we propose an approach that combines two complementary modeling steps. The first relies on generalized additive models (GAMs) to separate deterministic spatiotemporal effects from stochastic residuals. The second builds on nested vine copulas (NVCs), a hierarchical vine-merging framework that flexibly models joint distributions across variables and locations. The NVC construction combines two layers of dependence: (i) spatial dependence across locations for each variable, and (ii) inter-variable dependence at a selected reference location, represented through a multivariate vine structure linking the variable-specific spatial models into a coherent joint structure. The GAM layer ensures spatiotemporal consistency, while the NVC layer captures both spatial and inter-variable dependence. We show in the case study that our method matches or outperforms state-of-the-art approaches, with potential operational implications for future Swiss climate projections. The Swiss Federal Office of Meteorology and Climatology (MeteoSwiss) regularly provides up-to-date national climate scenarios that are freely available for a wide range of applications. In the two most recent releases, CH2018 (Fischer et al. 2022) and CH2025, bias correction is performed using empirical QM, which does not account for inter-variable, spatial, or temporal dependence. The methodology and results presented in this manuscript could therefore inform the next generation of Swiss climate scenarios and be directly relevant for future operational developments.

The remainder of the paper is organized as follows. [Sections 2](#) and [3](#) present the GAM and NVC layers, respectively. The full bias correction procedure is detailed in [Section 4](#). [Section 5](#) presents the application to our Swiss case study, and [Section 6](#) concludes with a discussion.

2 Generalized Additive Models for Bias Correction

A common strategy to address spatiotemporal dependence is to decompose the data into a systematic component and a stochastic remainder. In this context, generalized additive

models (GAMs, see e.g., [Hastie & Tibshirani 1990](#), [Wood 2025](#)) allow for nonlinear effects of covariates meant to capture the spatiotemporal structure such as time, latitude, longitude, altitude, and interactions thereof.

Let Y be a climate variable of interest, i.e., the bias correction target, and $\mathbf{X} \in \mathbb{R}^p$ be a vector of covariates. GAMs extend linear models by allowing the conditional mean of the response $\mu(\mathbf{x}) = \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}]$ depend on the covariates through a sum of additive components

$$g(\mu(\mathbf{x})) = \sum_{j=1}^m f_j(\mathbf{x}),$$

the m additive terms being linear or smooth functions, e.g., defined via spline bases, and $g(\cdot)$ is a link function connecting the linear predictor to the response. Inference is carried out by assuming that the response $Y \mid \mathbf{X} = \mathbf{x} \sim F_{\mu(\mathbf{x})}$, where F_{μ} denotes an exponential-family distribution with mean $\mu(\mathbf{x})$, leaving the dispersion parameter aside for simplicity. This offers flexible modeling while maintaining the additive interpretability of predictor effects.

From a fitted model, one can then compute the probability integral transform (PIT) $U = F_{\hat{\mu}(\mathbf{x})}(Y)$ to map the response to a uniform distribution on $[0, 1]$. For zero-inflated variables such as precipitation, a randomized PIT draws U uniformly on $[0, F_{\hat{\mu}(\mathbf{x})}(0)]$ when $Y = 0$. A GAM-based equivalent to the quantile mapping approach described in [Section 1](#) can then be written as

$$\tilde{Y} = F_{\hat{\mu}^{(rc)}(\mathbf{x})}^{-1} \circ F_{\hat{\mu}^{(mc)}(\mathbf{x})}(Y),$$

using the fitted GAMs for the reference and model data during the calibration period, respectively. To account for model misspecification, one can then add another QM-like layer in the middle of the transformation.

However, even when applied to a single location and a single variable, this approach has two drawbacks. First, it relies on GAM-based extrapolation to the projection period, i.e., the

absolute time is needed as a covariate, instead of, say, time-of-year, time-of-day, or other covariates strictly meant to capture a “deterministic” temporal structure. Second, and more importantly, it does not explicitly account for the modeled climate change signal. Our GAM-equivalent of the quantile delta mapping of Cannon et al. (2015) can be written as

$$\tilde{Y} = F_{\hat{\mu}^{(rc)}(\mathbf{x}) + \hat{\mu}^{(mp)}(\mathbf{x}) - \hat{\mu}^{(mc)}(\mathbf{x})}^{-1} \circ F_{\hat{\mu}^{(mp)}(\mathbf{x})}(Y), \quad (1)$$

where the model data during the projection period is transformed to the uniform scale using the GAM fitted on the same period, thus avoiding extrapolation in time issues; the location parameter in the inverse CDF is adjusted by the difference between the GAMs fitted on the projection and calibration periods, to preserve the modeled climate change signal.

This approach can be applied independently to each variable, and form the basis for our method. By construction however, it does not explicitly account for inter-variable dependencies. When considering a single location, one could add a vine bias correction (Funk et al. 2025) layer in the middle of the transformation. In the next section, we propose a class of models that extends this idea to the multi location setting.

3 Nested Vine Copulas

Copulas allow for flexible construction of a joint distribution with arbitrary one-dimensional margins (see e.g., Nelsen 2007, Joe 2014, for textbook treatments). Mathematically, a copula is a multivariate distribution with standard uniform margins. From Sklar (1959), for any multivariate distribution F with univariate marginal distributions F_i for $i = 1, \dots, d$, there exists a copula C such that $F(y_1, \dots, y_d) = C\{F_1(y_1), \dots, F_d(y_d)\}$ for all $y \in \mathbb{R}^d$, which is unique if the margins are continuous. It implies that the joint density is the product between the marginal densities and the copula density $c = \partial^d C / \partial u_1 \cdots \partial u_d$, so that the joint log-likelihood is the sum of the marginal log-likelihoods and that of the copula. This can be exploited in a two-step procedure by first estimating each of the margins separately

and then the copula (Genest et al. 1995, Joe & Xu 1996).

3.1 Vine Copulas

While standard copulas are well understood and have found many applications, their flexibility to model complex data structures is limited. This is because, in most copulas, dependencies between all subsets of variables are described by the same parametrization, which is often too restrictive. Popularized in Aas et al. (2009), vine copulas allow for a finer-grained modeling approach, and offer closed-form expressions for the required conditional distributions. Following the seminal work of Joe (1996) and Bedford & Cooke (2001, 2002), any d -variate copula density c can be decomposed into a product of $d(d-1)/2$ bivariate (conditional) copula densities. The order of conditioning in this decomposition can be organized using a graphical structure, called regular vine (R-vine) - a sequence of spanning trees $\mathcal{V} = \{\mathcal{T}_j\}_{j=1}^{d-1}$, with $\mathcal{T}_j = (N_j, E_j)$ the tree, nodes and edges at level j . The trees are nested in a way such that the nodes in the first tree correspond to the variables themselves, and the edges in a tree \mathcal{T}_j become the nodes of the next tree \mathcal{T}_{j+1} . Such a construction is formalized in the following definition.

Definition 3.1 (Czado (2019)) *A sequence of $d-1$ spanning trees is an R-vine if $N_1 = \{1, \dots, d\}$, $N_j = E_{j-1}$ for $j = 2, \dots, d-1$, and the **proximity condition** holds, i.e., if two nodes in \mathcal{T}_{j+1} are joined by an edge, the corresponding edges in \mathcal{T}_j share a common node.*

A vine copula identifies each edge of the sequence, its structure, with a bivariate copula density, called a pair-copula, and the density then factorizes as a product of these pair-copulas and the marginal densities.

Figure 3 shows an example with three variables. In tree \mathcal{T}_1 , variables 1 and 2 as well as variables 2 and 3 are linked by a pair-copula, the edge representing their connection serves as node in tree \mathcal{T}_2 , and




Figure 3: A 3d vine.

the corresponding copula density can be written as

$$c(u_1, u_2, u_3, u_4) = c_{12}(u_1, u_2) \cdot c_{23}(u_2, u_3) \cdot c_{34}(u_3, u_4) \cdot c_{13|2}(C_{1|2}(u_1|u_2), C_{3|2}(u_3|u_2)),$$

where $C_{1|2}$ and $C_{3|2}$ denote the conditional distribution functions of U_1 and U_3 given U_2 , respectively.

A vine-based multivariate distribution on d variables can then be represented by the triplet

$$F := \left(\{F_j\}_{j=1}^d, \mathcal{V}, \mathcal{C}(\mathcal{V}) \right), \quad (2)$$

where $\{F_j\}_{j=1}^d$ denotes the marginal distributions, \mathcal{V} a structure satisfying [Definition 3.1](#), and $\mathcal{C}(\mathcal{V})$ is the corresponding collection of pair-copulas. Because each margin and pair-copula can be specified independently, vine copula models are highly flexible and can be tailored to the specific needs of the application. We refer to recent books and surveys ([Joe 2014](#), [Aas 2016](#), [Czado 2019](#), [Czado & Nagler 2022](#)) for comprehensive overviews.

3.2 Nested Vine Copulas

In the spatial context, it is common to have separate models for different locations, and the question arises of how to combine them into a single model that captures the joint distribution across all locations. Or one could envision a situation where one has separate models for different variables, and wants to combine them into a single model that captures the joint distribution across all variables. In this section, we propose a new construction, nested vine copula (NVC), which provides a framework for combining two separate vine-based multivariate distributions in a hierarchical manner. While we illustrate NVC for two vine copulas, the concept naturally generalizes to $d \geq 2$ multivariate distributions.

Letting $\mathbf{Y}^{(1)} \in \mathbb{R}^{d_1}$ and $\mathbf{Y}^{(2)} \in \mathbb{R}^{d_2}$ be two (possibly dependent) random vectors, where d_1 and d_2 may differ, suppose that they follow a vine-based distribution as in [Equation \(2\)](#), denoted $F^{(1)}$ and $F^{(2)}$, respectively. NVC derives another vine-based distribution F of

$\mathbf{Y} = (\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)})^\top \in \mathbb{R}^{d_1+d_2}$ that is coherent with $F^{(1)}$ and $F^{(2)}$ by merging their graph structures and associated pair-copulas and capturing both shared and distinct dependence patterns. A hierarchical construction from separate vines has advantages over a new model for the pooled data. When related datasets are available (e.g., from different regions, time periods, or experimental settings), separate models can uncover context-specific dependence structures. Additionally, from a computational standpoint, merging existing vine structures reduces computational cost by reusing previously estimated components.

Since both $\mathbf{Y}^{(1)}$ and $\mathbf{Y}^{(2)}$ admit vine-based representations, the marginals of F are naturally defined as the union of the marginals of $F^{(1)}$ and $F^{(2)}$. To combine $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(2)}$ into a valid vine structure \mathcal{V} , we construct \mathcal{T}_j for $j = 1, \dots, d_1 + d_2 - 1$ as follows:

- (1) Add all nodes and edges from $\mathcal{T}_j^{(1)}$ and $\mathcal{T}_j^{(2)}$, that is $N_j^{(1)} \cup N_j^{(2)}$ and $E_j^{(1)} \cup E_j^{(2)}$, where $N_j^{(1)} = E_{j-1}^{(1)} = \emptyset$ for $j > d_2$ if $d_2 < d_1$ and conversely if $d_1 < d_2$.
- (2) Add a set of bridging edges $E_j^{(b)}$ satisfying the proximity condition to the disconnected graph $(N_j^{(1)} \cup N_j^{(2)} \cup E_{j-1}^{(b)}, E_j^{(1)} \cup E_j^{(2)})$, so that the resulting graph $(N_j^{(1)} \cup N_j^{(2)} \cup E_{j-1}^{(b)}, E_j^{(1)} \cup E_j^{(2)} \cup E_j^{(b)})$ is a spanning tree, starting with $E_0^{(b)} \equiv \emptyset$.

The merged vine structure \mathcal{V} thus combines the original dependence structures with additional cross-dependencies induced by the bridging edges. And we refer to Section A in the supplementary material for a theorem on the validity of this procedure. The associated set of pair-copulas is given by $\mathcal{C}(\mathcal{V}) = \mathcal{C}(\mathcal{V}^{(1)}) \cup \mathcal{C}(\mathcal{V}^{(2)}) \cup \mathcal{C}(\mathcal{V}^{(b)})$, where $\mathcal{C}(\mathcal{V}^{(b)})$ denotes the pair-copulas corresponding to the bridging edges. These additional copulas capture the dependence between nodes belonging to $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(2)}$, ensuring that the resulting model defines a coherent joint distribution.

[Figure 4a](#) illustrates this process. To initialize the recursive construction, a bridging edge $e \in \{j, k\}$ with $j \in \{1, 2, 3\}$ and $k \in \{4, 5\}$ has to be selected to connect the




Figure 4: Merging two vines with node sets $N_1^{(1)} = \{1, 2, 3\}$ and $N_1^{(2)} = \{4, 5\}$ with their edges as solid black lines and the bridging edges as solid blue lines. Panel (a) shows all trees with $\{3, 4\}$ as first tree bridging pick. Panel (b) shows the first two trees with $\{2, 4\}$ instead, with the dotted blue lines showing the two potential candidates to connect \mathcal{T}_2 .

nodes $N_1 = \{1, \dots, 5\}$ of \mathcal{T}_1 . Picking $e = \{3, 4\}$ arbitrarily, the merged \mathcal{T}_1 has edge set $E_1^{(1)} \cup E_1^{(2)} \cup \{e\} = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}\}$. In \mathcal{T}_2 , e becomes part of the node set and appears in both bridging edges. With this pick, only one configuration satisfies the proximity requirement, and the merging is uniquely determined afterwards. Choosing $e = \{2, 4\}$ instead, as shown in Figure 4b, the number of admissible bridging edges in \mathcal{T}_2 increases and one needs to further pick either $\{1, 4; 2\}$ or $\{3, 4; 2\}$, in addition to $\{2, 5; 4\}$, to form $E_2^{(b)}$. For $j \geq \max(d_1, d_2) = 3$, the original vines have empty edge sets, so the merging procedure proceeds solely by adding bridging edges until $\mathcal{T}_{d_1+d_2-1} = \mathcal{T}_4$ is obtained.

We give details about the implementation in Appendix A.1; we outline the main ideas here. Both steps operate on tree-edge lists rather than the usual matrix representations of vine structures (Czado 2019), which avoids redundant re-encoding and keeps the merge independent of any particular array format. Step (1) is implemented as concatenation of edge lists: for each level j , the edges from $\mathcal{V}^{(1)}$ and $\mathcal{V}^{(2)}$ are merged into $E_j = E_j^{(1)} \cup E_j^{(2)}$. Step (2) proceeds level by level. At each level, admissible bridging edges are first enumerated by checking the proximity condition. They are then ordered arbitrarily, manually, at random,

or according to an optimality criterion, analogously to the vine selection algorithm of [Dissmann et al. \(2013\)](#). Finally, they are processed greedily in that order in a Kruskal-like fashion ([Kruskal 1956](#)), with acyclicity enforced via a disjoint-set union structure ([Galler & Fischer 1964](#), [Tarjan 1975](#)), until a spanning tree is reached. The resulting merged vine is generally not unique and depends on the candidate-ordering rule. With $n_j = d + 1 - j$ with $d = d_1 + d_2$ the number of nodes at level j , candidate generation scans $O(n_j^2)$ node pairs, each requiring set operations on supports of size j , giving a per-level bound of $O(n_j^2 j)$. Summing across levels gives an overall worst-case complexity bound that simplifies to $O(d^4)$. In the structured setting of the next section however, the partially merged trees are nearly complete at low levels, and the routine is fast in practice for the dimensions considered.

4 Spatiotemporal Multivariate Bias Correction

In this section, we describe **GN-VBC**, our spatiotemporal multivariate bias correction approach, incorporating **GAMs** and **NVC** in the MBC method **VBC** ([Funk et al. 2025](#)). Our procedure uses three conceptual steps:

1. **Model**: the multivariate distributions of the reference data and the climate simulations, respectively in the calibration and projection periods.
2. **Align**: the climate simulations with the distribution of the reference data in the calibration period using the estimated distributions.
3. **Adjust**: the aligned data with the climate change signal, namely the discrepancy between the simulated evolution from the calibration to the projection period.

Our GN-VBC method is implemented in the R package **GNVBC**, available on GitHub ([Meier & Vatter 2026](#)). In Section B of the supplementary material, we compare GN-VBC to both QM and the VBC approach of [Funk et al. \(2025\)](#), illustrating how our method preserves spatiotemporal and multivariate relationships in a synthetic data setting. The remainder of

this section details each of the three steps.

4.1 Modeling Step

For d variables and s locations, let $Y^{(mp)} = \{Y_{i,j,t}^{(m)} : i = 1, \dots, d, j = 1, \dots, s, t \in p\}$ be the modeled projections, where i , j , and t represent respectively the variable, location and time indices, and similarly for the modeled and reference data in the calibration period. Treating each climate variable and dataset separately, we follow [Section 2](#) and fit a spatiotemporal GAM using a tensor-product spline in time and location, expressed as

$$g_i(\mu_{i,j,t}) = f_i(d_t, \text{Lat}_j, \text{Lon}_j),$$

where $\mu_{i,j,t} = \mathbb{E}[Y_i | \mathbf{X} = (t, j)]$, d_t is the day-of-year at timestamp t , Lat_j and Lon_j are the latitude and longitude at location j , and we dropped the superscript for the sake of clarity. For each dataset and variable, we then use the estimates to compute the PITs

$$U_{i,j,t} = F_i(Y_{i,j,t}; \hat{\mu}_{i,j,t}), \quad (3)$$

where $F_i(\cdot)$ denotes the cumulative distribution function of the fitted family for variable i . For each timestamp t , the PITs are then rearranged so that the values over all variables and locations form a single ds -dimensional vector

$$\mathbf{U}_t = (U_{1,1,t}, \dots, U_{1,s,t}, U_{2,1,t}, \dots, U_{d,s,t}) \in [0, 1]^{ds}.$$

Under the above ordering, the pair (i, j) corresponds to column $(i-1)s + j$ of \mathbf{U}_t . Extending [Funk et al. 2025](#) to the spatiotemporal context, one could then fit ds -dimensional unstructured vines separately to the samples $U^{(mp)} = \{\mathbf{U}_t^{(m)} : t \in p\}$ and $U^{(rc)} = \{\mathbf{U}_t^{(r)} : t \in c\}$. Instead, we exploit the NVC of [Section 3](#) and fit, for each PIT dataset,

- a location-specific vine on the PITs with indices $\{(i, j)\}_{j=1}^s$ for each variable i ,
- a variable-specific vine on the PITs with indices $\{(i, j_0)\}_{i=1}^d$ for a given j_0 ,

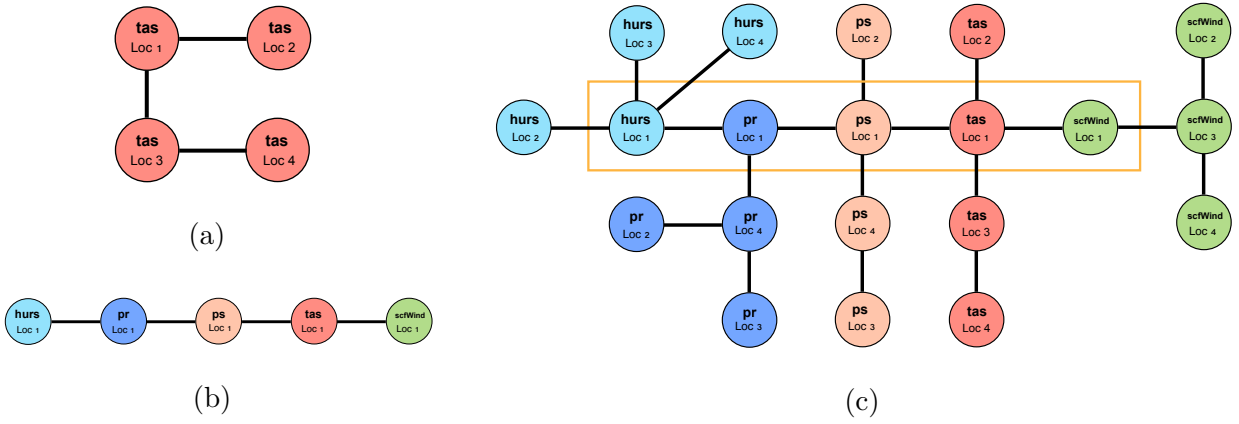


Figure 5: Example of merging location- and variable-specific vine structures based on four locations and five variables in a 20-dimensional vine copula. Panel (a) shows the first tree of the location-specific vine for temperature; Panel (b) the first tree of the variable-specific vine at location 1 which is highlighted in yellow in Panel (c) which shows the resulting merged first tree. Variable abbreviations are given in [Table 1](#).

where j_0 is a carefully chosen bridging location. For each PIT dataset, the d location-specific vines are then merged with the NVC procedure into ds -dimensional vine-based distributions $\hat{F}^{(mp)}$ and $\hat{F}^{(rc)}$, with the variable-specific vine providing the bridging edges that determine how variables are connected across locations. The two types of vines are complementary: the location-specific vines model within-variable spatial structure, while the variable-specific vine captures inter-variable dependence. An illustrative example is given in [Figure 5](#).

4.2 Alignment Step

As in [Funk et al. \(2025\)](#), we use the Rosenblatt transform ([Rosenblatt 1952](#)) and its inverse as multivariate analogues of the PIT and the quantile function to bias-correct $U^{(mp)}$. For a distribution F on \mathbb{R}^d , the transform and its inverse are maps $\mathcal{R}_F : \mathbb{R}^d \rightarrow [0, 1]^d$ and $\mathcal{R}_F^{-1} : [0, 1]^d \rightarrow \mathbb{R}^d$ defined recursively from a sequence of conditional distributions associated with F . If \mathbf{Y} is a random vector with distribution F , then $\mathcal{R}_F(\mathbf{Y})$ is a random vector with independent uniform components. Conversely, \mathcal{R}_F^{-1} maps a random vector with independent uniform components into one that has distribution F . Although \mathcal{R}_F and \mathcal{R}_F^{-1} are generally

unavailable in closed form, vine-based distributions are a notable exception.

It allows us to use the NVC-based ds -dimensional estimated distributions $\hat{F}^{(mp)}$ and $\hat{F}^{(rc)}$ to transfer the dependence structure of the reference calibration period to the biased model projections by successive applications of the Rosenblatt transform and its inverse, that is

$$\tilde{\mathbf{U}}_t^{(m)} = \mathcal{R}_{\hat{F}^{(rc)}}^{-1} \circ \mathcal{R}_{\hat{F}^{(mp)}} \left(\mathbf{U}_t^{(m)} \right), t \in p,$$

a multivariate analogue to the quantile mapping described in [Section 1](#). Using the projection-period vine copula model $\hat{F}^{(mp)}$ rather than $\hat{F}^{(mc)}$ avoids imposing stationarity of the dependence structure across periods and allows the bias correction to preserve the multivariate rank dynamics projected by the climate model.

4.3 Adjustment Step

To map the bias-corrected PITs $\tilde{U}^{(mp)} = \{ \tilde{\mathbf{U}}_t^{(m)} : t \in p \}$ back to bias-corrected climate projections, we proceed as in [Equation \(1\)](#) in [Section 2](#) and define

$$\tilde{Y}_{i,j,t}^{(m)} = F_{\hat{\mu}_{i,j,t}^{(rc)} + \hat{\mu}_{i,j,t}^{(mp)} - \hat{\mu}_{i,j,t}^{(mc)}}^{-1} \left(\tilde{U}_{i,j,t}^{(m)} \right),$$

where $\hat{\mu}_{i,j,t}^{(rc)}$, $\hat{\mu}_{i,j,t}^{(mc)}$, and $\hat{\mu}_{i,j,t}^{(mp)}$ denote the fitted GAM means for the reference calibration, model calibration, and model projection datasets, respectively. This reverses the PIT step in [Equation \(3\)](#) while preserving the modeled climate change signal through the additive correction $\hat{\mu}_{i,j,t}^{(mp)} - \hat{\mu}_{i,j,t}^{(mc)}$ in the location parameter.

5 Application

We apply the proposed methodology to the dataset introduced in [Section 5.1](#). The aim is to bias-correct daily simulations of five atmospheric variables ([Table 1](#)) at 22 grid points in the canton of Vaud (Switzerland). The calibration period is 1980–2009 and the projection period 2010–2022. Results compare the proposed GN-VBC pipeline to a set of benchmark methods summarized in [Table 2](#).

Method	Uni.	Multi.	Indep.	Site-wise	Joint	Raw	PIT	NVC
QM	×		×			×		
MBCn		×			×	×		
R2D2		×			×	×		
VBC		×		×		×		
C-VBC		×			×	×		
G-VBC		×		×			×	
GC-VBC		×			×		×	
N-VBC		×			×	×		×
GN-VBC		×			×		×	×

Table 2: Overview of bias correction methods. Columns 2 and 3 indicate whether the method is univariate or multivariate. Columns 4, 5 and 6 show how the method is applied: independently to each location and variable, site-wise, or jointly to all location-variable pairs. Columns 7 and 8 indicate whether raw data or PIT residuals are used within the method, and column 9 whether our vine merging algorithm, NVC, is incorporated.

As a univariate baseline, we use empirical QM. The multivariate methods considered fall into two categories: methods applied separately at each of the 22 locations (VBC and G-VBC) and methods applied jointly across all locations (all remaining methods). In the first category, VBC is used either directly on the raw time series or on the probability integral transform (PIT) values obtained from GAM decompositions (G-VBC) based on a smooth effect for the day of the year only. The second category includes two state-of-the-art MBC algorithms: MBCn proposed by Cannon (2018), as well as R2D2 proposed by Vrac (2018). In their work, Cannon (2018) apply MBCn site-wise, while we follow François et al. (2020) who showcase its application to the joint dataset. For this study, we compared both performances (not shown) and opted for the joint application due to better results and, therefore, a stronger comparison with our proposed methods. For both MBCn and R2D2

we use the implementation in the R package `MBC`. Note that for `R2D2`, the package requires the same projection and reference period lengths, meaning that the calibration period narrows down to 1980–1992. In addition, we consider two naïve joint approaches based on standard VBC: C-VBC, where VBC is applied to all variable-location pairs jointly using the raw time series, and GC-VBC, where the raw time series is replaced by PIT residuals. These approaches are considered naïve because the Dißmann algorithm (as implemented in `rvinecopulib` (Nagler & Vatter 2025)) which is used in the implementation of VBC) is free to connect arbitrary variables across space, potentially leading to unintuitive pairings (e.g., humidity at one location paired with wind speed at a distant location).

This limitation motivates the N-VBC setting, where our NVC merging procedure constrains the vine structure to follow interpretable spatial and inter-variable relationships. Finally, GN-VBC denotes our entire proposed MBC pipeline as introduced in Section 4, which applies our NVC merging algorithm to PIT values obtained from GAM decompositions based on a spatiotemporal tensor spline to estimate coherent nested vine structures across all variables and locations. This sequential setup isolates the contributions of the spatiotemporal GAM decomposition and the NVC merging step, allowing us to assess their effects independently. Note that in both G-VBC and GC-VBC, the delta mapping of VBC is replaced by our version of adjusting the mean structure (see Section 4.3) to not account for varying climate conditions on the PIT scale.

More details on the fit of the nested vine copula model of the entire GN-VBC pipeline are provided in the supplementary material. Next to an evaluation of our marginal model fit in Section C, we show the first trees of the fitted nested vine models in Section D, providing a graphical representation of the dependencies modeled by GN-VBC.

The choice of GAM family in step 4.1 depends on the distributional characteristics of each variable. In this case study, across all methods using GAMs, we model temperature

(**tas**) and surface pressure (**ps**) with a Gaussian distribution, wind speed (**sfcWind**) with a Gamma distribution, relative humidity (**hurs**) with a Beta distribution due to its support on $[0, 1]$, and precipitation (**pr**) with a Tweedie distribution to account for the zero-inflated mixed discrete–continuous nature of rainfall.

The NVC merging uses a bridging location chosen near the canton centre ([Figure 1b](#)) to provide a stable anchor. To preserve geographic interpretability we optionally restrict the first tree of variable-specific vines to edges between adjacent grid points. Joint variable–location vines are truncated at level 22 for computational tractability; alternative truncation depths produced only minor changes in evaluation metrics.

Uncertainty is quantified with an annual block bootstrap ([Künsch 1989](#)): entire years are resampled with replacement (block size 20) jointly across all variables and locations, thereby preserving within-year temporal and spatial dependence.

5.1 Data

We focus on five atmospheric variables (see [Table 1](#)) across the canton of Vaud in Switzerland (see [Figure 1](#)) for the period 1 January 1980 to 31 December 2022, at daily temporal resolution and 0.11° spatial resolution. Our reference data combine observational and reanalysis products. Specifically, temperature and precipitation are obtained from the TabsD and RhiresD grid-data products by MeteoSwiss ([MeteoSwiss 2021](#)). Relative humidity, wind speed, and surface pressure are derived from ERA5-Land (ECMWF Reanalysis 5th Generation for Land; see [Hersbach et al. 2025](#)) by averaging the hourly values.

Climate simulations are sourced from the Coordinated Regional Climate Downscaling Experiment for the European domain (EURO-CORDEX) at the resolution of 0.11° (EUR-11), using the RCP8.5 scenario ([Giorgi et al. 2009](#), [Jones et al. 2011](#)). The simulations are driven by the CNRM-CERFACS-CM5 global model and downscaled with the CLMcom-

ETH-COSMO-crCLIM-v1-1 regional model over the EUR-11 domain. For consistency, both the MeteoSwiss observational datasets and ERA5-Land fields are remapped onto the common CORDEX grid using bilinear and conservative remapping, respectively (Mueller & Schulzweida 2010). The period 1980–2009 is used for calibration, and 2010–2022 for evaluation against the reference data.

5.2 Inter-variable Dependencies

As in François et al. (2020) and Funk et al. (2025), we assess the preservation of inter-variable dependencies by comparing the reference and the bias-corrected distributions at each location using the Wasserstein distance (Villani 2008). For two distributions F_1 and F_2 , its square is given by $W_2^2(F_1, F_2) = \inf_{\gamma \in \Pi(F_1, F_2)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\gamma(x, y)$, with $\Pi(F_1, F_2)$ the set of distributions on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals F_1 and F_2 , and $\|\cdot\|$ the Euclidean norm. After computing 5-dimensional second Wasserstein distances over our climate variables for each location, we show in Figure 6 the relative improvement

$$\frac{W_2(F^{(rp)}, F^{(mp)}) - W_2(F^{(rp)}, \tilde{F}^{(mp)})}{W_2(F^{(rp)}, F^{(mp)})}, \quad (4)$$

where $\tilde{F}^{(mp)}$ corresponds to the distribution of the bias-corrected model projections. Across all methods, a broadly consistent spatial pattern emerges, with substantial improvements typically between 70 and 90 % for most grid points, and comparatively lower improvements in the northeastern part of the canton. Despite being a univariate method that does not explicitly model inter-variable dependence, QM achieves improvements comparable to several multivariate approaches in this setting.

Among the multivariate methods, R2D2, MBCn, and VBC yield similar levels of improvement. Incorporating the GAM-based decomposition prior to VBC (G-VBC) does not lead to an additional gain in this purely location-wise setting. In contrast, when locations are treated jointly, C-VBC, without structural constraints on the dependence across variables and locations, results in the lowest overall improvement. This performance loss is par-




Figure 6: Improvement in multivariate second Wasserstein distances for five climate variables per location and BC method. The mean value over all locations is indicated in the title.

tially mitigated by incorporating the GAM decomposition in GC-VBC and more effectively addressed by introducing nested vines in N-VBC, which enforces spatially coherent and interpretable vine structures. Combining both the GAM-based decomposition and the NVC merging algorithm yields improvements comparable to those achieved by established multivariate methods. Overall, these results indicate that our proposed approach matches existing MBC techniques in terms of inter-variable consistency. The following subsections assess its ability to further preserve spatial and temporal consistency.




Figure 7: Evaluation of spatial consistency: improvement in second Wasserstein distance with respect to no correction (top row, higher is better); MSE between the spatial correlation of the reference and the respective correction method (bottom row, lower is better).

5.3 Spatial Consistency

In our proposed framework, spatial coherence is promoted at two stages: first through the spatiotemporal decomposition via location-specific tensor splines, and second through the hierarchical merging of variable-specific vine copulas using NVC. To quantitatively assess how well spatial dependence is preserved, we follow [François et al. \(2020\)](#) and evaluate all methods using variable-specific second Wasserstein distances. Relative improvements with respect to the reference are computed as in [Equation \(4\)](#), but based on the joint 22-dimensional distributions across all locations. The first row of [Figure 7](#) shows the resulting improvements for three selected variables as boxplots over bootstrap replicates; numerical

values for all variables are provided in Section G of the supplementary material.

For near-surface temperature (left column), median improvements range between approximately 26% and 72% for most methods, with the notable exception of VBC applied independently at each location. Additional analyses (not shown) indicate that this poor performance is primarily driven by strong seasonal patterns in temperature. Since standard VBC operates directly on the raw time series and does not explicitly account for seasonality, vine estimation becomes unstable in this setting. This effect largely disappears when applying C-VBC, where the joint modeling of all variables and locations allows the vine copula to borrow strength across dimensions. Introducing the GAM-based decomposition prior to bias correction leads to a substantial improvement for both location-wise and joint applications of VBC (G-VBC and GC-VBC). Moreover, imposing a hierarchical structure via NVC further improves spatial consistency compared to unrestricted joint modeling. The strongest gains are achieved when combining GAM decomposition with NVC in GN-VBC. This suggests a complementary interaction between the two: the GAM decomposition stabilizes marginal behavior by removing systematic spatiotemporal structure, while NVC enforces an interpretable spatial dependence structure. R2D2, which has been shown to preserve spatial dependencies across multiple studies ([François et al. 2020](#), [Vrac 2018](#)), performs comparably well to our approach, GN-VBC. As for MBCn, although it is applied to all locations simultaneously, it performs poorly.

A similar pattern is observed for relative humidity (right column), although improvements are generally more moderate than for temperature. Applying VBC independently at each location yields only limited improvements. Incorporating the GAM decomposition (G-VBC) leads to a noticeable enhancement, but the overall performance remains modest. This indicates that removing seasonal and smooth spatiotemporal effects alone is insufficient to recover realistic spatial dependence in relative humidity. A substantial improvement is

achieved when VBC is applied jointly to all locations as in C-VBC. Moreover, individually adding the GAM decomposition or the NVC merging strategy to the joint VBC setup does not lead to marked additional gains. However, when both components are combined in GN-VBC, just as discussed for temperature, a clear improvement emerges, with similar performance as R2D2. In contrast, empirical QM and MBCn show little to no improvement for relative humidity, reflecting its inability to address spatial dependence structures.

For precipitation (middle column), the differences between the correction methods are less pronounced, with median improvements ranging between approximately 16% and 27%. Incorporating GAM decompositions (G-VBC and GC-VBC) slightly deteriorates the performance of both VBC and C-VBC. This suggests that, for precipitation, removing smooth spatiotemporal components is less beneficial. Introducing NVC leads to a modest improvement in median performance over C-VBC, indicating that enforcing a structured and interpretable spatial dependence is advantageous even when overall gains are limited. As observed for the other variables, the combination of GAM decomposition and NVC (GN-VBC) ultimately outperforms competing approaches, including R2D2, demonstrating that jointly addressing marginal stabilization and spatial dependence remains beneficial.

While Wasserstein distances quantify similarity between joint spatial distributions, a complementary indicator of spatial consistency is the preservation of spatial correlation between locations. Following [Largeau et al. \(2025\)](#), we define

$$B_r = \left\{ (s_1, s_2) \in \{1, \dots, s\}^2 : d(s_1, s_2) \in \left[r - \frac{\Delta r}{2}, r + \frac{\Delta r}{2} \right] \right\}, \quad d_r = |B_r|.$$

as the set of all pairs of locations whose pairwise distances fall within a distance bin centered at r with radius Δr . For standardized (in time) $\tilde{Y}_{i,j,t}$ of length N , i.e., $t = 1, \dots, N$, the spatial correlation function for variable i is defined by

$$\rho_{\tilde{Y}_i}(r) = \frac{1}{N} \sum_{t=1}^N \frac{1}{d_r} \sum_{(s_1, s_2) \in B_r} (\tilde{Y}_{i,s_1,t} - \langle \tilde{Y}_{i,t} \rangle) (\tilde{Y}_{i,s_2,t} - \langle \tilde{Y}_{i,t} \rangle) / \text{std}_{\tilde{Y}_i},$$

where for each time t we define the spatial mean within the bin as

$$\langle \tilde{Y}_{i,t} \rangle = \frac{1}{d} \sum_{j=1}^d \tilde{Y}_{i,j,t} \quad \text{and} \quad \text{std}_{\tilde{Y}_i} = \frac{1}{d} \sum_{j=1}^d \left(\tilde{Y}_{i,j,t} - \langle \tilde{Y}_{i,t} \rangle \right)^2.$$

In our setting, the distance bins are chosen to correspond to the length of the shortest path between location pairs when only adjacent grid cells are allowed to connect. This results in six distinct radii, with the largest radius representing pairs of locations connected through six adjacent edges.

For each variable and BC method, the spatial correlation function is evaluated at these six radii and compared to the reference spatial correlation function using the MSE. The results across the 20 bootstrap replicates are shown as boxplots in the second row of [Figure 7](#). For temperature, the raw model output already reproduces the spatial correlation structure reasonably well, leaving limited room for improvement. Consequently, most BC methods yield comparable performance, with the exception of VBC and G-VBC, which exhibit higher variability in MSE values. For precipitation and relative humidity, the variability across bootstrap samples is substantially larger, reflecting the more heterogeneous spatial structure of these variables. Among all methods, GN-VBC shows the biggest variance reduction across bootstrap replicates, indicating a more stable preservation of spatial correlation. In the case of relative humidity, VBC-based methods that model all locations jointly generally outperform location-wise corrections. In particular, C-VBC provides a clear improvement over standard VBC, highlighting the importance of explicitly modeling spatial dependence. Interestingly, the GAM decomposition alone does not directly enhance the preservation of spatial correlations. However, when combined with the hierarchical vine structure imposed by NVC, it leads to similar performance as R2D2. This suggests that the benefits of removing spatiotemporal trends via GAMs are fully realized only when coupled with a structured multivariate dependence model. Neither univariate MBCn nor QM improves upon the uncorrected model output for either precipitation or relative humidity. This





Figure 8: Spatial autocorrelation as a function over shortest path lengths between grid points for three variables. The dashed black line shows the reference. For clarity, only R2D2, and the proposed GN-VBC approach are shown.

confirms that marginal corrections alone are insufficient to maintain spatial coherence.


To further illustrate the behavior of the spatial correlation function across the six distance classes considered in this study, Figure 8 shows that, as expected, across all variables spatial correlation generally decreases with increasing path length. For larger path lengths (five to six adjacency edges), a slight increase in correlation is observed. This effect arises from the elongated geometry of the canton of Vaud: due to its skewed shape (see Figure 1b), some location pairs that are geographically close in Euclidean distance are connected only through longer adjacency paths, leading to higher correlations at larger path lengths than for some intermediate distances.

While the differences between the model and the reference and bias-corrected temperature data are marginal, the differences between the model and the reference precipitation and relative humidity data are more pronounced at larger distances. Both R2D2 and GN-VBC perform equally well in improving the spatial correlation of the model, demonstrating that our novel MBC approach can compete with existing spatially coherent methods.

Overall, these results confirm that the hierarchical vine structure imposed by NVC, in combination with the marginal stabilization provided by the GAM decomposition, leads to a more faithful reproduction of spatial dependence patterns. More evaluations regarding



(a)



(b)

Figure 9: Lag- and location-wise squared difference in ACF values with respect to reference projections per variable and method for 365 days (Panel a) and 30 days (Panel b). Grey lines show squared differences per location; smoothed lines show the mean over locations. inter-variable - spatial dependencies are provided in Section E of the supplementary material.

5.4 Temporal Consistency

Finally, we assess the proposed BC framework in terms of temporal consistency. Since the GAM decomposition explicitly models smooth seasonal and spatiotemporal effects through the day-of-year component, methods incorporating this step are expected to better preserve temporal dependence structures by construction.

To quantify temporal consistency, we evaluate the autocorrelation function (ACF) at lags $t = 1, \dots, 365$ for each variable-location pair. For each lag, we compute the squared

difference between the ACF of the bias-corrected series and that of the reference projection, yielding a lag-specific measure of temporal distortion; see [Figure 9a](#) for lags up to one year and [Figure 9b](#) for short-term dependence up to 30 days. Summaries for all variables are in the supplementary material (Section G).

For near-surface temperature, the raw climate model output already reproduces the temporal autocorrelation structure of the reference reasonably well, leaving little scope for improvement, and QM, MBCn, and R2D2 all show similar performance. VBC applied independently at each location shows some degradation of temporal correlations at individual sites, but its performance remains acceptable. In contrast, methods that apply VBC jointly across locations without accounting for seasonality (C-VBC and N-VBC) distort the ACF, with pronounced peaks at seasonal lags, around two weeks, six months, and one year. The GAM decomposition effectively removes these artifacts: GAM-based variants (G-VBC, GC-VBC and GN-VBC) closely align with the reference ACF across all lags.

For precipitation, temporal differences are generally smaller across all methods. The uncorrected model output already shows low ACF deviations, and both QM and MBCn preserve this behavior, while R2D2 struggles to preserve short-term dependencies ([Figure 9b](#)). Similarly, VBC-based approaches applied to all locations combined without GAM decomposition (C-VBC and N-VBC) show medium discrepancies at very short lags, suggesting slight distortion of short-term temporal consistency. Incorporating the GAM decomposition removes these early-lag discrepancies. Overall, all methods perform comparably well for precipitation in terms of temporal consistency.

Relative humidity presents the most challenging case. The raw model output clearly deviates from the reference ACF, exhibiting a seasonal pattern similar to that observed for temperature under joint VBC corrections (C-VBC), though less pronounced. Similar patterns are shown by QM and location-based VBC, while MBCn partially improves them without fully

eliminating seasonal artifacts. Again, R2D2 distort short-term temporal dependencies and exhibits similar seasonal patterns as MBCn. VBC-based methods that ignore seasonality lead to pronounced temporal inconsistencies by amplifying these deviations. In contrast, all methods incorporating the GAM decomposition yield substantial improvements. Among these, the proposed GN-VBC approach achieves the closest agreement with the reference ACF, indicating a synergistic effect between seasonal adjustment and structured multivariate dependence modeling.

Although R2D2 preserves inter-variable and spatial correlations well, it clearly struggles to maintain temporal consistency. This is addressed by our GN-VBC approach. Our results demonstrate that incorporating a GAM-based decomposition prior to bias-correcting is effective in preserving temporal dependence structures. While multivariate dependence modeling alone is insufficient to ensure temporal consistency, combining GAMs with the hierarchical dependence structure enforced by NVC yields robust performance across variables.

6 Discussion

With the aim of advancing operational climate projections for Switzerland (an inherently challenging task given the country’s complex topography) we propose a new multivariate bias correction method that captures inter-variable and spatiotemporal dependence more explicitly and flexibly than existing approaches, and demonstrate its strong performance on our case study. The approach builds on generalized additive models (GAMs) and nested vine copulas (NVCs). In the current implementation, location-specific vine copulas are merged through a variable-based bridging vine defined at a selected spatial location. Preliminary experiments with the inverse merging strategy, i.e., constructing per-location vines across variables and merging them through a spatial bridging vine, performed less favorably. This likely reflects the fact that individual climate variables typically exhibit stronger spatial dependence than multiple variables observed at a single location. A systematic comparison

of alternative merging strategies remains a promising avenue for future research.

Several extensions of the framework are possible. The present implementation is tailored to gridded climate data and would require additional development to accommodate irregularly spaced locations (e.g., [Worku et al. 2020](#)). In addition, the bridging location is currently chosen heuristically; evaluating multiple candidates and selecting the optimal one according to objective criteria may further improve performance.

The GAM decomposition used in this study models smooth spatiotemporal variations in the mean but does not explicitly target tail behavior. Understanding how the proposed method affects extremes, particularly under nonstationary dependence structures, therefore remains an important direction for future work. Incorporating additional covariates, such as altitude, into the GAM formulation may also improve the representation of spatial variability.

Although motivated by climate bias correction, the NVC framework is broadly applicable to multivariate data with hierarchical dependence structures. Potential applications include joint modeling of outcomes across multiple hospitals in clinical studies or financial risk assessment across countries (e.g., [Allen et al. 2017](#)). Overall, the results illustrate the potential of hierarchical vine constructions as a flexible and interpretable approach for modeling complex multivariate dependence structures.

7 Disclosure statement

The authors report there are no competing interests to declare.

8 Data Availability Statement

Both [CORDEX](#) and [ERA5-Land](#) data are publicly available in the Climate Data Store. MeteoSwiss TabsD and RhiresD [grid-data products](#) are publicly available. Their remapped version to the EUR-11 0.11° spatial resolution may be provided on request.

References

- Aas, K. (2016), ‘Pair-copula constructions for financial applications: A review’, *Econometrics* **4**(4), 43.
- Aas, K., Czado, C., Frigessi, A. & Bakken, H. (2009), ‘Pair-copula constructions of multiple dependence’, *Insurance: Mathematics and Economics* **44**(2), 182–198.
- Allen, D. E., McAleer, M. & Singh, A. K. (2017), ‘Risk measurement and risk modelling using applications of vine copulas’, *Sustainability* **9**(10), 1762.
- Bárdossy, A. & Pegram, G. (2012), ‘Multiscale spatial recorrelation of RCM precipitation to produce unbiased climate change scenarios over large areas and small’, *Water Resources Research* **48**(9).
- Bedford, T. & Cooke, R. M. (2001), ‘Probability density decomposition for conditionally dependent random variables modeled by vines’, *Annals of Mathematics and Artificial Intelligence* **32**(1-4), 245–268.
- Bedford, T. & Cooke, R. M. (2002), ‘Vines—a new graphical model for dependent random variables’, *The Annals of Statistics* **30**(4), 1031–1068.
- Cannon, A. J. (2016), ‘Multivariate bias correction of climate model output: Matching marginal distributions and intervariable dependence structure’, *Journal of Climate* **29**(19), 7045–7064.
- Cannon, A. J. (2018), ‘Multivariate quantile mapping bias correction: an n-dimensional probability density function transform for climate model simulations of multiple variables’, *Climate Dynamics* **50**(1), 31–49.
- Cannon, A. J., Sobie, S. R. & Murdock, T. Q. (2015), ‘Bias correction of GCM precipitation by quantile mapping: how well do methods preserve changes in quantiles and extremes?’, *Journal of Climate* **28**(17), 6938–6959.
- Czado, C. (2019), ‘Analyzing dependent data with vine copulas’, *Lecture Notes in Statistics*,

Springer **222**.

- Czado, C. & Nagler, T. (2022), ‘Vine copula based modeling’, *Annual Review of Statistics and Its Application* **9**(1), 453–477.
- Dissmann, J., Brechmann, E. C., Czado, C. & Kurowicka, D. (2013), ‘Selecting and estimating regular vine copulae and application to financial returns’, *Computational Statistics & Data Analysis* **59**, 52–69.
- Ehret, U., Zehe, E., Wulfmeyer, V., Warrach-Sagi, K. & Liebert, J. (2012), ‘HESS Opinions “Should we apply bias correction to global and regional climate model data?”’, *Hydrology and Earth System Sciences* **16**(9), 3391–3404.
- Erhardt, T. M., Czado, C. & Schepsmeier, U. (2015a), ‘R-vine models for spatial time series with an application to daily mean temperature’, *Biometrics* **71**(2), 323–332.
- Erhardt, T. M., Czado, C. & Schepsmeier, U. (2015b), ‘Spatial composite likelihood inference using local C-vines’, *Journal of Multivariate Analysis* **138**, 74–88.
- Fischer, A. et al. (2022), ‘Climate scenarios for Switzerland CH2018 – Approach and implications’, *Climate Services* **26**, 100288.
- Flato, G. et al. (2013), Evaluation of climate models, in ‘Climate Change 2013: The Physical Science Basis.’, Cambridge University Press, pp. 741–866.
- François, B., Vrac, M., Cannon, A. J., Robin, Y. & Allard, D. (2020), ‘Multivariate bias corrections of climate simulations: which benefits for which losses?’, *Earth System Dynamics* **11**(2), 537–562.
- Funk, H., Ludwig, R., Küchenhoff, H. & Nagler, T. (2025), ‘Towards more realistic climate model outputs: a multivariate bias correction based on zero-inflated vine copulas’, *Journal of the Royal Statistical Society Series C: Applied Statistics* p. qlaf044.
- Galler, B. A. & Fischer, M. J. (1964), ‘An improved equivalence algorithm’, *Communications of the ACM* **7**(5), 301–303.
- Genest, C., Ghoudi, K. & Rivest, L.-P. (1995), ‘A semiparametric estimation procedure of

- dependence parameters in multivariate families of distributions’, *Biometrika* **82**(3), 543–552.
- Giorgi, F., Jones, C., Asrar, G. R. et al. (2009), ‘Addressing climate information needs at the regional level: the CORDEX framework’, *World Meteorological Organization (WMO) Bulletin* **58**(3), 175.
- Gräler, B. (2014), ‘Modelling skewed spatial random fields through the spatial vine copula’, *Spatial Statistics* **10**, 87–102.
- Hastie, T. & Tibshirani, R. (1990), *Generalized Additive Models*, Chapman and Hall.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I. et al. (2025), ‘ERA5 monthly averaged data on single levels from 1979 to present’, *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)* **10**, 252–266.
- Ivanov, M. A. & Kotlarski, S. (2017), ‘Assessing distribution-based climate model bias correction methods over an alpine domain: added value and limitations’, *International Journal of Climatology* **37**(5), 2633–2653.
- Joe, H. (1996), Families of m -variate distributions with given margins and $m(m - 1)/2$ bivariate dependence parameters, in ‘Distributions with fixed marginals and related topics’, IMS Lecture Notes - Monograph Series, pp. 120–141.
- Joe, H. (2014), *Dependence Modeling with Copulas*, CRC Press.
- Joe, H. & Xu, J. J. (1996), ‘The estimation method of inference functions for margins for multivariate models’, *Technical Report (Department of Statistics, University of British Columbia)* **166**.
- Jones, C., Giorgi, F. & Asrar, G. (2011), ‘The Coordinated Regional Downscaling Experiment: CORDEX, an international downscaling link to CMIP5’, *CLIVAR Exchanges* **56**.
- Kotlarski, S., Keuler, K., Christensen, O. B., Colette, A., Déqué, M., Gobiet, A., Goergen,

- K., Jacob, D., Lüthi, D., Van Meijgaard, E. et al. (2014), ‘Regional climate modeling on European scales: a joint standard evaluation of the EURO-CORDEX RCM ensemble’, *Geoscientific Model Development* **7**(4), 1297–1333.
- Kruskal, J. B. (1956), ‘On the shortest spanning subtree of a graph and the traveling salesman problem’, *Proceedings of the American Mathematical Society* **7**(1), 48–50.
- Künsch, H. R. (1989), ‘The jackknife and the bootstrap for general stationary observations’, *The Annals of Statistics* **17**(3), 1217–1241.
- Largeau, L., Beucler, T., Leutwyler, D., Mariethoz, G., Chavez-Demoulin, V. & Koch, E. (2025), ‘Investigating the robustness of extreme precipitation super-resolution across climates’, *Weather and Climate Extremes* (**to appear**).
- Meier, T. & Vatter, T. (2026), *GN-VBC: Nested Vine Copula based Multivariate Bias Correction for Climate Projections*. R package version 0.0.0.9000.
URL: <https://github.com/TheresaMeier/GNVBC>
- MeteoSwiss (2021), Meteoswiss spatial climate analyses: documentation of datasets for users, Technical report, Federal Office of Meteorology and Climatology MeteoSwiss.
- Mueller, R. & Schulzweida, U. (2010), ‘Climate data operators’, *Version* **1**(5.1), 614.
- Nagler, T. & Vatter, T. (2025), *rvinecopulib: High Performance Algorithms for Vine Copula Modeling*. R package version 0.7.3.1.0.
- Nelsen, R. B. (2007), *An Introduction to Copulas*, Springer Science & Business Media.
- Rajczak, J., Kotlarski, S., Salzmann, N. & Schaer, C. (2016), ‘Robust climate scenarios for sites with sparse observations: a two-step bias correction approach’, *International Journal of Climatology* **36**, 1226.
- Robin, Y., Vrac, M., Naveau, P. & Yiou, P. (2019), ‘Multivariate stochastic bias corrections with optimal transport’, *Hydrology and Earth System Sciences* **23**(2), 773–786.
- Rosenblatt, M. (1952), ‘Remarks on a multivariate transformation’, *The annals of mathematical statistics* **23**(3), 470–472.

- Rummukainen, M. (2016), ‘Added value in regional climate modeling’, *WIREs Climate Change* **7**(1), 145–159.
- Sklar, A. (1959), ‘Fonctions de répartition à n dimensions et leurs marges’, *Publications de L’Institut de Statistique de L’Université de Paris* **8**, 229–231.
- Tarjan, R. E. (1975), ‘Efficiency of a good but not linear set union algorithm’, *Journal of the ACM* **22**(2), 215–225.
- Vautard, R., Kadygrov, N., Iles, C., Boberg, F., Buonomo, E., Bülow, K., Coppola, E., Corre, L., van Meijgaard, E., Nogherotto, R. et al. (2021), ‘Evaluation of the large EURO-CORDEX regional climate model ensemble’, *Journal of Geophysical Research: Atmospheres* **126**(17), e2019JD032344.
- Villani, C. (2008), *Optimal Transport: Old and New*, Springer.
- Vrac, M. (2018), ‘Multivariate bias adjustment of high-dimensional climate simulations: the Rank Resampling for Distributions and Dependences (R2D2) bias correction’, *Hydrology and Earth System Sciences* **22**(6), 3175–3196.
- Wood, S. N. (2025), ‘Generalized additive models’, *Annual Review of Statistics and Its Application* **12**(1), 497–526.
- Worku, G., Teferi, E., Bantider, A. & Dile, Y. T. (2020), ‘Statistical bias correction of regional climate model simulations for climate change projection in the Jemma sub-basin, upper Blue Nile Basin of Ethiopia’, *Theoretical and Applied Climatology* **139**, 1569–1588.