



HAL
open science

A temporal stochastic bias correction using a machine learning attention model

Omer Nivron, Damon Wischik, Mathieu Vrac, Emily Shuckburgh, Alex Archibald

► To cite this version:

Omer Nivron, Damon Wischik, Mathieu Vrac, Emily Shuckburgh, Alex Archibald. A temporal stochastic bias correction using a machine learning attention model. *Environmental Data Science*, 2025, 3, pp.e36. <10.1017/eds.2024.42>. <hal-04893289>

HAL Id: hal-04893289

<https://hal.science/hal-04893289v1>

Submitted on 17 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.


L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

METHODS PAPER

A temporal stochastic bias correction using a machine learning attention model

Omer Nivron¹ , Damon J. Wischik¹, Mathieu Vrac², Emily Shuckburgh¹ and Alex T. Archibald³

¹Department of Computer Science and Technology, University of Cambridge, Cambridge, UK

²Laboratoire des Sciences du Climat et de l'Environnement, Institut Pierre-Simon Laplace, Paris, France

³Yusuf Hamied Department of Chemistry, University of Cambridge, Cambridge, UK

Corresponding author: Omer Nivron; Email: on234@cam.ac.uk

Received: 27 June 2024; **Accepted:** 28 September 2024

Keywords: attention; bias correction; heatwaves; machine learning; time series

Abstract

Climate models are biased with respect to real-world observations. They usually need to be adjusted before being used in impact studies. The suite of statistical methods that enable such adjustments is called bias correction (BC). However, BC methods currently struggle to adjust temporal biases. Because they mostly disregard the dependence between consecutive time points. As a result, climate statistics with long-range temporal properties, such as the number of heatwaves and their frequency, cannot be corrected accurately. This makes it more difficult to produce reliable impact studies on such climate statistics. This article offers a novel BC methodology to correct temporal biases. This is made possible by rethinking the philosophy behind BC. We will introduce BC as a time-indexed regression task with stochastic outputs. Rethinking BC enables us to adapt state-of-the-art machine learning (ML) attention models and thereby learn different types of biases, including temporal asynchronicities. With a case study of number of heatwaves in Abuja, Nigeria and Tokyo, Japan, we show more accurate results than current climate model outputs and alternative BC methods.

Impact Statement

Climate models are biased. Bias correction (BC) is a suite of statistical methods to address those biases. Corrected climate model statistics enable more precise impact studies. However, to date, BC struggles to adjust for long-range temporal properties such as the number of heatwaves. We offer a novel methodology for stochastic temporal BC. This is possible by rethinking BC as a time-indexed regression model with stochastic outputs. This reformulation enables using state-of-the-art machine learning (ML) attention models which can learn different types of biases, including temporal asynchronicities. We show more accurate results than current climate model outputs and alternative BC methods on the number of heatwaves statistics case study in Abuja, Nigeria and Tokyo, Japan. The results are relevant to a broad audience, ranging from climate scientists to policymakers.

1. Introduction

Climate models are biased with respect to observations (Palmer & Weisheimer, 2011; Christensen et al., 2008; Teutschbein et al., 2012), affecting impact studies in domains such as the economy (Richard, 2018; OECD, 2015; Richard, 2021), human migration (Defrance et al., 2017), biodiversity (Bellard et al., 2012), hydrology (Gleick, 1989; Christensen et al., 2004; Piao et al., 2010), and agronomy (Ciais et al., 2005;

Ari, 2020). If left uncorrected, those biases can invalidate the outcomes of impact studies (e.g., for hydrological model (Grouillet et al., 2015)). (See Stainforth et al. (2007) for a detailed discussion on reasons for these biases.)

Climate scientists have developed a suite of tools termed model output statistics (MOS) methods (defined in Maraun & Widmann, 2018, p. 5) that correct these biases and provide output statistics that more closely match when compared to observations. Bias correction (BC) is a subset of MOS methods. However, the term BC is used differently in different communities (Maraun & Widmann, 2018, p. 171). We use BC as synonymous with homogeneous (same physical variable) MOS with or without a spatial resolution mismatch.

Most BC methods are univariate, correcting for a single one-dimensional physical variable in a single location, such as the summer daily mean air temperatures in London (Ho et al., 2012). The most popular univariate methods are mean correction (Xu, 1999), variance correction (Boberg & Christensen, 2012; Ho et al., 2012; O'Reilly et al., 2020; Hawkins et al., 2013; Iizumi et al., 2011; Eden et al., 2012; Schmidli et al., 2006), and empirical quantile mapping (EQM) along with its variations (Ho et al., 2012; Li H & Wood, 2010; Salvi et al., 2021; Wood et al., 2002; Hagemann et al., 2011; Themesl et al., 2011; Turco et al., 2017; Déqué, 2007; Vrac et al., 2012; Kallache et al., 2010). Additionally, there are multivariate BC approaches. They correct properties of multidimensional distributions (Robin et al., 2019; Piani & Haerter, 2012; Cannon, 2018; Dekens et al., 2017; Vrac, 2018)) such as the joint distribution of daily temperatures and rainfall at multiple locations.

However, BC methods currently struggle to adjust for temporal biases. As Vrac and Thao (2020) explain, univariate methods cannot correct biases relating to temporal dependencies. Because they discard temporal information (see a simple illustration of the argument in Figure 31). This issue can be exacerbated by a spatial resolution mismatch between the reference and climate model data since univariate BC methods will struggle to reproduce the grid-scale temporal variability. For example, in the EQM method, a future projection value will be large if the climate model value is large, regardless of the specific observed behavior at the corresponding grid point. Additionally, most multivariate BC approaches are not designed to correct for statistics relying on the dependence between consecutive time points. Thus, they also struggle to correct temporal biases, as pointed out by Vrac and Thao (2020) and shown by François et al. (2020).

There have been attempts to create BC methods that explicitly capture temporal properties. Vrac and Friederichs (2015) developed a method (“EC-BC”) where the ranks of the climate simulations are shuffled to match those of observations. However, this method implies that the ranks of a future time series are the same as those in the present period. In this study, we show that its estimation of number of heatwaves can be inaccurate when the assumption does not hold. Another method, by Mehrotra and Sharma (2019) (“3D-BC”), extends the Vrac et al. method to include 1-lag auto-correlation. This method has the advantage that it can capture the correlation of two consecutive readings by construction. However, extending it to higher-order correlations is not trivial. Vrac and Thao (2020) introduced a rank resampling-based BC method (“R2-D2v2”), which aims to adjust multivariate properties while constraining the temporal properties to be similar to those of the references. However, it cannot be used in a homogeneous MOS setup since it requires an additional pivot variable assumed to be correctly ranked. More recently, Robin and Vrac (2021) proposed the “TSMBC” method, which relies on using the multivariate BC method “dOTC” (Robin et al., 2019) in a particular configuration. It enables the adjustment of temporal dependencies by incorporating lagged time series as additional variables to be corrected. However, the authors show that limiting the temporal constraints to only a few timestamps is required. Consequently, it is limited in estimating different statistics, such as long heatwaves. Despite these gaps, known impact studies about future heatwaves rely on univariate BC methods or the raw outputs of climate models (Marx et al., 2021; Summers et al., 2022).

The goal of this paper is to introduce a BC method that can correct climate models and more accurately estimate future climatic statistics with long-term temporal properties, such as heatwaves.

We find it helpful to offer a new BC perspective to achieve this goal. We cast the classic BC—typically treated as an algorithmic procedure mapping between histograms—as a probability model with a specific

time index and without independence assumptions. Having a time-indexed target random variable as part of the regression task does not mean we can precisely specify the value obtained in Tokyo, Japan, on January 1st, 2050. However, there are still multiple benefits, such as (a) we can obtain the distribution at that specific time point, and this distribution can represent the typical behavior for that month, season or decade and (b) we can account for temporal correlations—what happened in day t can affect the possible values at day $t+1$. (For a detailed discussion on motivation for our probabilistic perspective, see [Appendix B](#) and associated [Figures 16–18](#).)

We are aware that our approach is unconventional, mainly since climate models and observations are not in temporal synchrony. However, by including past observations and (past, future) climate model outputs and choosing a flexible ML architecture, we can learn the distribution at a specific time point and temporal correlations despite the temporal asynchrony. To implement our BC method, we have devised a novel ML methodology based on the Taylorformer attention model (Nivron et al., 2023; Vaswani et al., 2017).

To put our method in context, we note that other methods model temporal properties explicitly. However, these methods are not strictly homogeneous MOS methods like ours. These methods include some form of perfect prognosis, i.e., they link statistically large-scale phenomena with local-scale phenomena in the observed data and apply the learned link to the climate model. Hence, we do not detail them further or use them as benchmarks. (For an example, see a conditional weather generator by Chandler (Chandler, 2020).)

To showcase the potential of our methodology, we test it on two case studies. We aim to correct the number of heatwaves statistics in Abuja, Nigeria and Tokyo, Japan, which have very different climate characteristics and distribution shifts (see [Figures 6–15](#), [Appendix A](#)). We use the simplified number of heatwaves definition for our case studies (as a proof of concept): at least three consecutive days above a chosen absolute value in $^{\circ}\text{C}$. Our method performs accurately on the standard BC evaluations, such as QQ plots. Further, our model consistently outperforms the alternatives on the number of heatwaves statistic for different thresholds during 20 years (1989–2008). For the 22°C threshold in Tokyo, our model misestimated the observed heatwaves by just 1% (averaged over 32 initial-condition runs). In contrast, the next-best BC model, “TSMBC,” misestimated the observed heatwaves by 8% (averaged over 32 initial-condition runs). For the 24°C threshold in Abuja, our model misestimated the observed heatwaves by 0.5%, whereas the second-best BC model, in this case, the mean shift, misestimated the observed heatwaves by 9%. Lastly, our model outperforms all the other methods on the log-likelihood score. We believe these results support a further extension of this study to check its generalization potential. (See results in the experiment Section 4 and extended results in [Appendix F](#) and associated [Figures 25–30](#).)

2. Reference and model data

In any BC task, data from a climate model are matched to real-world observations (or a proxy of observations, namely a reanalysis product) to correct biases. Because Global Climate Models (GCMs) are typically coarser in resolution than observations or reanalysis reference data, the BC task is often joined by a downscaling of resolution.

We follow the same approach. For the climate model data, we chose the 32 initial condition runs from the climate model of the Institut Pierre-Simon Laplace (IPSL) (Lurton et al., 2020) (Boucher et al., 2020) (Hourdin et al., 2020) under the sixth International Coupled Model Intercomparison Project (CMIP6) (Eyring et al., 2016) historical experiment. The data are obtained at a 1-day frequency for the variable “tasmax,” representing the maximum temperature 2 meters above the surface. The IPSL model is run at a 250-km nominal resolution and is not regridded. For the case studies, we take the closest geographical point to Abuja, Nigeria, and the closest point to Tokyo, Japan, from 1948 to 2008.

For the observational reference data, we use two different Reanalysis datasets:

- For Abuja, Nigeria, we use the NCEP-NCAR Reanalysis 1 (Kalnay et al., 1996) of the National Oceanic and Atmospheric Administration (NOAA). The data are obtained in six-hourly daily

measurements for the variable “tmax,” the maximum temperature 2 meters above the surface, and are averaged to daily simulations. The resolution is 2.5 degrees globally in both the north–south and east–west directions and is not regridded. We chose the closest geographical point to Abuja, Nigeria.

- For Tokyo, Japan, we use the ERA5 reanalysis (Hersbach et al., 2020) from the European Centre for Medium-Range Weather Forecasts (ECMWF). The data are obtained daily for the variable “mx2t,” the maximum temperature 2 meters above the surface. The resolution is 0.25 degrees globally in both the north–south and east–west directions and is not regridded. We chose the closest geographical point to Tokyo, Japan.

We use two different reanalysis datasets, NCEP-NCAR and ERA5, to show that the methodology is not limited to one spatial resolution configuration. We hypothesize that the model will also work for many different configurations of (climate model, reanalysis) pairs.

3. Methodology

To start, let us first describe the big picture of our method. Our goal in this study is to estimate potential continuations of observed values from 1989 to 2008, given past observations and (past, future) climate model outputs. To do so, we (1) generate a custom dataset comprised multiple (climate model, observation) pairs with varying time-series lengths from the reference and climate model data, (2) train a probability model with an explicit likelihood function¹ on the custom data, and (3) generate entire time-series samples from the trained probability model to continue the reference data to the future (we call this step the “Inference procedure”, as is common in the ML literature). We use the generated samples to estimate different statistics. Here, we concentrate on estimating the distribution of “number of heatwaves.”

In the following, we will describe our approach in greater detail, starting in reverse order to explain the motivation of our end goal. We first detail the generation process using a trained probability model (step 3) before explaining how to train it (steps 1 and 2). To be more precise, we introduce mathematical notation. Throughout the text, we will use uppercase letters to denote random variables and lowercase letters to denote specific values. Let O_i be a random variable representing real-world reading of maximum daily temperatures at time t_i , and g_i , the GCM output of maximum daily temperatures at time t_i . Bolded variables are vectors, and the notation $i:j$ denotes the set of all integer indices from h through j , inclusive.

3.1. Inference procedure

Using the notation above, our goal (step 3 above) is to estimate $\mathbf{O}^* = (O_1^*, \dots, O_m^*)$ at some time points t_1^*, \dots, t_m^* given GCM outputs $\mathbf{g} = (g_1, \dots, g_\ell)$ at time points t'_1, \dots, t'_ℓ and observed real-world readings $\mathbf{o} = (o_1, \dots, o_n)$ at t_1, \dots, t_n . We do not assume synchronicity: the climate model time-points t'_1, \dots, t'_ℓ may in general be different to times $t_1, \dots, t_n, t_1^*, \dots, t_m^*$.

Since we start in reverse order, we assume we already have trained some probabilistic model that can generate samples of \mathbf{O}^* given \mathbf{g} and \mathbf{o} (and time-points). We will additionally suppose that this probability model admits an explicit likelihood function, call it $P_\theta(\mathbf{O}^*|\mathbf{g}, \mathbf{o})$, where θ represents the learnable parameters. We call \mathbf{o} and \mathbf{g} the conditioning set.

Generating samples of \mathbf{O}^* given \mathbf{g} and \mathbf{o} is a useful goal for estimating many different statistics. Because, for example, if we sample many time series, we can estimate the marginal distribution of a specific time point, such as O_1^* . Separately, we can estimate the average of \mathbf{O}^* . Furthermore, for each generated time series, we can estimate the number of consecutive days above a chosen threshold (“number of heatwaves”).

¹ A likelihood function is a measure of how likely a given set of parameters would result in the observed data. Some models, like hidden Markov models, do not have a closed-form likelihood function.




Figure 1. Inference Task for Estimating Future Observations: The top panel outlines the available data to us in order to estimate the potential continuation of observed values post-1989 (dashed vertical line), based on historical observations (green line) and both past and future climate model outputs (blue line). The bottom panel displays a red line representing one possible continuation, sampled from the forecasting model, which estimates $P_{\theta}(\mathbf{O}^*|\mathbf{o},\mathbf{g})$.

Figure 1 shows in the top panel the data available to us at the start of the inference task, and the bottom panel shows an example of a generated sample time-series.

However, training a probability model to generate sensible time series—ones that yield accurate statistics compared to observed hold-out set data—is difficult because climate models and observations are not synchronized temporally (Maraun & Widmann, 2018), and we do not know the exact form of asynchrony.

The following section provides the blueprint for training a chosen probability model. After which, in Section 3.3, we discuss the specific architecture choice we have made in this study.

3.2. Training procedure

Assume we start with a training dataset that consists of real-world readings o_1, \dots, o_n at t_1, \dots, t_n and climate model outputs g_1, \dots, g_ℓ at t'_1, \dots, t'_ℓ . We emphasize in the notation that the climate model time points do not have to be synchronized with the real-world reading time points. However, we will assume they are synchronized for the readability of the data generation and training algorithms.

As described in Section 3.1, our goal is to train a probability model that can generate time-series samples. One such probability model is the Taylorformer (Nivron et al., 2023), described in Section 3.3, which we use for our evaluation.

In Algorithm 1, we describe the precise steps to generate a training batch. The motivation for the procedure will become apparent in Section 3.3. In loose language, an element i in the training batch is a pair of (climate model, observations) time series. The element is constructed by (i) randomly picking a window of 60–360 aligned (in-time) points both from observations and GCM outputs. We denote the resulting climate model time series as \mathbf{g}^i . Then (ii) randomly picking a time point t_j . This time point divides the observed time series, such that all points in the observations after t_j become the predictands, \mathbf{O}^{*i} and all points before t_j are given in \mathbf{o}^i . The dataset creation procedure is schematically shown in Figure 1. The dashed orange lines represent the chosen window, while the black vertical line shows t_j .

After constructing a training batch, we can use Algorithm 2 to optimize the parameters θ . Note that we have set the elements \mathbf{O}^{*i} , \mathbf{o}^i , and \mathbf{g}^i when constructing the batch in Algorithm 1 and we have assumed we have an explicit likelihood function. So, we just need to plug the elements into the

likelihood function and perform the optimization. (We call the likelihood score the “Training objective”).

Algorithm 1 Procedure to generate a random training batch.

- 1: **Input:** batch size B , observed sequence o_1, \dots, o_n , climate model sequence g_1, \dots, g_n (and associated time-points)
 - 2: **for** $i = 1$ **to** B **do**
 - 3:
 - “Time window boundary selection” for the training sample.
 - 4: Pick k uniformly at random from $\{1, \dots, n - 360\}$ ▸ left boundary
 - 5: Pick h uniformly at random from $\{k + 60, \dots, k + 360\}$ ▸ right boundary
 - 6:
 - “Prediction index selection” within the window.
 - 7: Pick j uniformly at random from $\{k + 5, \dots, h - 5\}$
 - 8:
 - Construct training sample i .
 - 9: Set $C_i = \{v | t_k \leq t_v \leq t_j\}$ ▸ Context indices
 - 10: Set $A_i = \{v | t_j < t_v \leq t_h\}$ ▸ Prediction indices
 - 11: Set $\mathbf{O}^{*,i} = [o_v | v \in A_i]$ ▸ Elements to be predicted
 - 12:
 - observations and climate model elements for conditioning set.
 - 13: Set $\mathbf{o}^i = [o_v | v \in C_i]$
 - 14: Set $\mathbf{g}^i = [g_v | v \in C_i \cup A_i]$
 - 15: **end for**
-

Algorithm 2 Training algorithm.

- 1: **while true do** ▸ break loop when objective is bigger than some chosen threshold
 - 2: Generate a random batch of time-series pairs of size B ▸ using Algorithm 1
 - 3: Perform gradient descent to maximize the objective:
 - 4: $\sum_{i=1}^B \log P_\theta(\mathbf{O}^{*,i} | \mathbf{g}^i, \mathbf{o}^i)$ ▸ “Training objective”
 - 5: **end while**
-

The probability model we use in Algorithm 2 can be theoretically any model that can generate samples from $P_\theta(\mathbf{O}^* | \mathbf{o}, \mathbf{g})$ has an explicit likelihood form and uses the constructed training data to do so. For example, a linear regression model can be used as the forecasting model, but it assumes by construction independence among the elements of \mathbf{O}^* , it will, most likely, produce a bad fit.

The inference procedure in Section 3.1 and the training procedure in Section 3.2 show the general purpose framework. We now discuss the specific choice of probability model we have made in this study.

3.3. Taylorformer as the probability model

In this study, we have chosen the Taylorformer (Nivron et al., 2023), which is a sequential ML model adapted from the Transformer attention model by Vaswani et al. (2017) to real-valued targets on a continuum. In Appendix D, we provide a detailed explanation of the Taylorformer. The explanation starts with the Transformer decoder for language modeling and the needed adaptations to reach the Taylorformer for BC. The motivation for choosing Taylorformer is its ability (not shown here) to learn different types of biases such as mean bias, variance inflation or deflation, and temporal a-synchronicity.

The Taylorformer learns these biases from data without specifying a priori the parametric form of the bias, which is exactly what we need since we do not know the time asynchronicity between GCMs and observations. Furthermore, the Taylorformer and other attention-based models can learn changing “rules” for temporal synchronicities. Their success in dealing with “temporal asynchronicity” is shown when translating a sentence from one language to another. For example, the subject of a sentence may change its position from one language to another. We now explain the adaptations to the Taylorformer needed for the BC task.

Training objective. Since the Taylorformer is a sequential model, it is a standard practice to decompose the “Training objective” in Algorithm 2 according to the probability chain rule. We write the training objective as

$$\begin{aligned} &\text{maximise}_{\theta} \log P_{\theta}(\mathbf{O}^{*,i} | \mathbf{o}^i, \mathbf{g}^i) = \\ &\text{maximise}_{\theta} \text{Log} P_{\theta}(O_h^i | O_{j:h-1}^i, \mathbf{o}^i, \mathbf{g}^i) \dots P_{\theta}(O_{j+1}^i | O_j^i, \mathbf{o}^i, \mathbf{g}^i) P_{\theta}(O_j^i | \mathbf{o}^i, \mathbf{g}^i) \end{aligned} \tag{1}$$

We model each term in the decomposition using a Normal distribution. Specifically, for any integer index s where s is bigger than the “prediction index” j , we define

$$O_s^i | \text{inp}_s^i \sim \text{Normal}(f_{\theta}(\text{inp}_s^i), g_{\theta}(\text{inp}_s^i)) \tag{2}$$

where we define $\text{inp}_s^i = (O_j^i, \dots, O_{s-1}^i, \mathbf{o}^i, \mathbf{g}^i)$. In our model, $f_{\theta}(\text{inp}_s^i)$ represents the mean value of the normally distributed O_s^i and $g_{\theta}(\text{inp}_s^i)$ its variance. Therefore, we get a distribution for O_s^i , i.e., stochastic outputs. Here, the functions f_{θ} and g_{θ} are defined by the Taylorformer architecture.

Training dataset. Practically, to train our dataset toward our inference goal, we must do more than feed in the entire time series of the pair (observations, climate model) from 1948 to 1988. Feeding in the entire time series would be fine for classical models such as ARIMA (G & Jenkins, 1970) or GPs. Differently, sequential deep learning models, such as the Taylorformer or a recurrent neural network (Rumelhart et al., 1986), cannot learn from one sequence since this would enable only one “gradient descent pass,” and little learning would occur. For that reason, in the batch generation Algorithm (1), we break down the pair of (climate model, observations) into multiple pairs, as shown in the first two columns of Figure 2.

We have to address two possible issues with our newly constructed sequences. (1) We want to allow pairs to have differing lengths, as it is unclear a priori which length is optimal; and (2) For a deep learning sequential model, we have to heed that it does not learn the rule Tomorrow will be the same as today as pointed out by Nivron et al. (2023). The Taylorformer can address both points if we feed it in with an appropriate dataset, as explained next.

To address point (1), the Taylorformer (Nivron et al., 2023) architecture can receive as input arbitrary length sequences by construction. Therefore, we simply need to feed in sequences of differing lengths as part of our training procedure (see “Time window selection” in Algorithm 1).

To address point (2), we employ the following “Pruning” mechanism: after we set $\mathbf{O}^{*,i}$ and $\mathbf{o}^i, \mathbf{g}^i$ in Algorithm 1, we randomly prune elements from them. Let us denote the results of the operation as $\text{Prune}(\mathbf{O}^{*,i})$, $\text{Prune}(\mathbf{o}^i)$, and $\text{Prune}(\mathbf{g}^i)$. We call the time points associated with $\text{Prune}(\mathbf{O}^{*,i})$ the “prediction tufts”. We show the “Pruning” operation and the “prediction tufts” in the third (rightmost) column of Figure 2. Intuitively, if the model does not receive as input the value at the previous time point at each iteration, it must search for a different “rule” for correction instead of copying yesterday’s value. Further, we limit the maximum sequence length to 360 points, which is chosen to address the high computational load of attention-based models such as the Taylorformer.

So we can now refine our training objective, incorporating the solutions to the two points above. The optimization becomes

$$\text{maximise}_{\theta} \text{Log} P_{\theta}(\text{Prune}(\mathbf{O}^{*,i}) | \text{Prune}(\mathbf{o}^i), \text{Prune}(\mathbf{g}^i)) \tag{3}$$

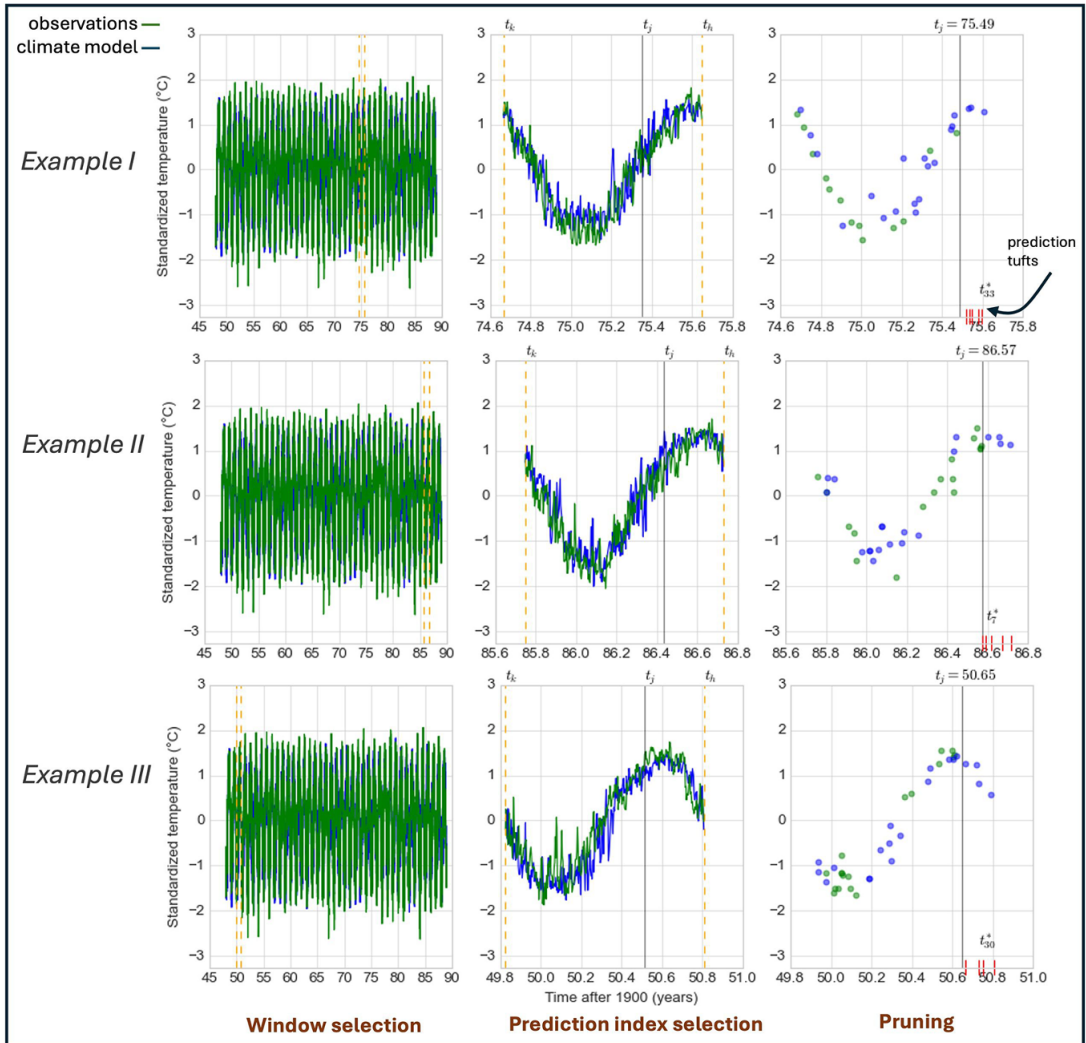


Figure 2. Construction of Training Data for Tokyo, Japan: Column I displays full sequences from 1948 to 1988 for climate models (blue) and observations (green). Dashed orange lines indicate the selected slice for each data row, a process termed “Window Selection.” Column II zooms into the selected window, featuring a vertical black line at a randomly selected “prediction index” (t_j), from which we aim to estimate the observations until t_h . Column III illustrates the “Pruning” operation, where time points before and after t_j are randomly selected from both the climate model and observations. The observed values to be estimated are concealed, and the chosen time points, called “prediction tufts” are highlighted with red tufts. This column adapts the data for use in ML sequential models. Note that we have dropped the example index i for readability, but k, j, h change from one example to the next.

which we again decompose according to the probability chain rule. (Note that the training dataset could be constructed in different ways. We discuss other possibilities in Appendix E.)

Multiple initial condition runs. Our source data consists of one climate model (“IPSL”) that was initialized 32 times to generate 32 different time-series. To include all the runs in our model, we add a “random run selection” step as the first command in the for loop in Algorithm 1.




Figure 3. Illustration of the inference procedure. Given a trained Taylorformer model with parameters θ , we generate a time-series sequentially. In the top row, we generate a value o_1^* which will be then plugged into the condition set in the second row and so forth.

Inference. In ML, we first train a model to learn the parameters of a network. Once the parameters are learned using an objective (here, maximum log-likelihood), they are fixed, and we can use the network to get our desired readouts. This is called the “Inference stage.” Our chosen Taylorformer model operates using sequential generation. In each step, we sample a value from a conditional distribution and plug the result into the condition set of the following conditional distribution. Figure 3 shows an illustration of the procedure. The bottom plot of Figure 1 shows an example of a fully sampled time series.

The inference differs from the training procedure in which we attach the observed value (as opposed to a sampled value) to the new condition set. This is the standard way to train a sequential model using maximum likelihood (also known as “teacher forcing” (Williams & Zipser, 1989)).

Due to computational constraints, we do not include in the conditioning set, \mathbf{o}, \mathbf{g} , all the observed data and (past, future) GCM data. Instead, we feed in the following sequence: \mathbf{o} contains the last 60 days of 1988. \mathbf{g} contains the last 60 days of 1988 and the next 120 days of 1989 from the GCM outputs. Then, given \mathbf{o} and \mathbf{g} , we produce the first estimated observation of 1989, O_1^* . Next, we attach the sampled value to our observations and shift all the processes by one day to the right so that when sampling the second day of 1989, O_2^* , we still have a condition set of the same size. We iterate the process until we have all days up to 2008. (Note that the inference procedure could be done in different ways. We discuss other possibilities in Appendix E.)

4. Experiments

To test our model, we do the following experiment: First, we bias-correct daily maximum temperatures for Abuja, Nigeria and separately for Tokyo, Japan, for 1989–2008 based on the years 1948–1988. Second, we calculate the number of heatwaves distribution, our primary interest, from the bias-corrected data.

A primary interest in using BC methods is to estimate the climate in the future. Here, we use 1989–2008 as our unseen “future.” Both Tokyo and Abuja have experienced a distribution shift (see Appendix A) in maximum daily temperatures between our training period, 1948–1988, and the unseen “future” period, 1989–2008, and hence serve as a good testing ground for our method.

An important question for BC methods is whether they preserve the trend exhibited by the climate model. Unlike mean shift, our method does not necessarily preserve the trend of the climate model per initial condition run. We show it in Appendix F in Figures 23 and 24 for Tokyo and Abuja, respectively. The figures represent the 5-year averages from 1989 to 2008 for different initial condition runs for the climate model (“GCM”), reanalysis data (“ERA”) and our Taylorformer model (“ML”). The figures also show that different initial condition runs from the climate model exhibit different trends, and in turn, those trends, in many cases, do not match the “observed” (reanalysis) trend.

Comparison. We compare our model to the raw output from the GCM (IPSL) and to six BC methods: mean correction (Xu, 1999), mean and variance correction (Ho et al., 2012), Empirical Quantile Mapping (EQM; Panofsky & Brier, 1968), empirical copula–bias correction (EC-BC; Vrac & Friederichs, 2015), three-dimensional bias correction (3DBC; Mehrotra & Sharma, 2019), and Time Shifted Multivariate Bias Correction (TSMBC; Robin & Vrac, 2021). Note that due to the spatial resolution mismatch, the mean correction, mean and variance correction, and EQM methods potentially misrepresent the local scale variability (see Maraun & Widmann, 2018, pp. 189–192). Nevertheless, we also compare with them since these are very popular BC methods (and, in turn, used for heatwave assessments) and are practically used with a spatial mismatch on many occasions, for example, in Vrac and Friederichs (Vrac & Friederichs, 2015). The full implementation details of the different BC methods are provided in Appendix I.

4.1. Results

Figure 4 shows the results for the “number of heatwaves” distribution for Tokyo, Japan using a threshold of 22°C. The raw outputs from the IPSL GCM (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) by more than 100% for the period 1989–2008. Our novel




Figure 4. Comparative Analysis of “number of heatwaves” Trends in Tokyo, Japan (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 22°C is shown. The IPSL climate model predictions are represented by red triangles, which generally underestimate the observations. Actual observations are indicated by a vertical orange line. The Taylorformer temporal BC are depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.




Figure 5. Comparative Analysis of “number of heatwaves” Trends in Abuja, Nigeria (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 24°C is shown. The IPSL climate model predictions are represented by red triangles, which overestimate the observations. A vertical orange line indicates actual observations. The Taylor former temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

Taylorformer temporal BC produces a much more accurate distribution (horizontal box plots) per GCM initial condition run (0–31) for the number of heatwaves in the same period. The Taylorformer’s average difference is 0.9% to the observed number of heatwaves. Other BC models perform worse: 17.5%, 8.5%, 3.1%, 22.7%, 47.5%, 1.1% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.

Figure 5 shows the results for the “number of heatwaves” distribution for Abuja, Nigeria, using a threshold of 24°C. The raw outputs from the IPSL GCM (red triangles) overestimate the number of instances of observed heatwaves (vertical orange line) by more than 33% for the period 1989 – 2008. Our novel Taylorformer temporal BC produces a much more accurate distribution (horizontal box-plots) per GCM initial condition run (0–31) for the number of heatwaves in the same period. The Taylorformer’s average difference is 0.5% to the observed number of heatwaves. Other BC models perform worse: 10%, 15%, 11%, 24%, 34%, 56% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.

For completeness, we provide QQ plots (Figures 19–20) and auto-correlation plots (Figures 21–22) in Appendix F. From the QQ plots, we see that our model performs at least as well as the EQM. The auto-correlation plots are more challenging to interpret. We also provide the MSE and log-likelihood scores in

Table 1. Our method is more accurate on both. The log-likelihood on a hold-out set gives us a combined measure of skill and spread and is thus informative. In contrast, the MSE is not very informative, typically, for evaluating BC tasks on the resulting time series since it does not provide us with a measure of the spread of the process. Nevertheless, since our method continues the observed time series, the MSE may be informative on short time scales (we leave the analysis of this point for future work).

5. Conclusions

This paper offers a novel stochastic temporal BC method. We test it on the correction of daily maximum temperatures time series in Abuja, Nigeria and Tokyo, Japan. It performs accurately on standard BC evaluations such as QQ plots and is also consistently more accurate than other BC methods on number of heatwaves statistics. The method is constructed by first rethinking BC as a time-indexed probability model and then adapting state-of-the-art attention ML models. Our model opens up the opportunity to do BC using a framework different from what is currently common, with potential accuracy gains regarding temporal properties and no apparent loss about the more usual distributional evaluations. Subsequently, these accuracy gains may have a large impact on policy-making decisions. Going forward, this study needs to be scaled up, considering multiple regions, more climate models, and more climate statistics. This can be done, for example, by applying it to experiments within dedicated frameworks, such as the VALUE framework (Maraun et al., 2015).

Open peer review. To view the open peer review materials for this article, please visit <http://doi.org/10.1017/eds.2024.42>.

Data availability statement. All data and code are open-sourced and available: IPSL climate model (data) – <https://esgf-node.llnl.gov/search/cmip6/>, NCEP-NCAR Reanalysis 1 (data) – <https://psl.noaa.gov/data/gridded/data.ncep.reanalysis.html>, ERA5 (data) <https://www.ecmwf.int/en/forecasts/dataset/ecmwf-reanalysis-v5>, Taylorformer temporal BC model (code) – <https://github.com/oremnirv/Taylorformer>

Author contribution. Omer Nivron: Conceptualization, Methodology, Writing—Original draft preparation, Investigation, Visualization, Software, Formal analysis. Damon Wischik: Conceptualization, Formal analysis, Methodology, Writing—Reviewing and Editing. Mathieu Vrac: Formal analysis, Writing—Reviewing and Editing, Validation, Supervision. Emily Shuckburgh: Writing—Reviewing and Editing, Supervision. Alexander Archibald: Writing—Reviewing and Editing, Supervision.

Provenance. This article was accepted into the Climate Informatics 2024 (CI2024) Conference. It has been published in Environmental Data Science on the strength of the CI2024 review process.

Funding statement. This research was supported by the Artificial Intelligence for Environmental Risks (AI4ER) CDT, University of Cambridge. Mathieu Vrac's work is also supported by state aid managed by the National Research Agency under France 2030 bearing the references ANR-22-EXTR-0005 (TRACCS-PC4-EXTENDING project) and ANR-22EXTR-0011 (TRACCS-PC10--LOCALISING project).

Competing interest. None.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Ari TB** (2020) Causes and Implications of the Unforeseen 2016 Extreme Yield Loss in France's Breadbasket. Available at <https://api.semanticscholar.org/CorpusID:235024720>.
- Céline Bellard**, et al. (2012) Impacts of climate change on the future of biodiversity. *Ecology Letters* 15(4), 365–377. doi: <https://doi.org/10.1111/j.1461-0248.2011.01736.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1461-0248.2011.01736.x>. url: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1461-0248.2011.01736.x>.
- Boberg F and Christensen J.** (2012) Overestimation of mediterranean summer temperature projections due to model deficiencies. *Nature Climate Change* 2, 433–436.
- Boucher O**, et al. (2020) Presentation and evaluation of the IPSL-CM6A-LR climate model. *Journal of Advances in Modeling Earth Systems* 12(7), e2019MS002010. doi: <https://doi.org/10.1029/2019MS002010>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS002010>. url: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS002010>.

- Box G EP and Jenkins GM** (1970) *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day
- Brown TB**, et al. (2020) Language Models are Few-Shot Learners. *ArXiv abs/2005.14165*. Available at <https://api.semanticscholar.org/CorpusID:218971783>.
- Camilo Mora** et al. (2017) Global risk of deadly heat. *Nature Climate Change* 7 doi: [10.1038/nclimate3322](https://doi.org/10.1038/nclimate3322).
- Cannon A** (2018) Multivariate quantile mapping bias correction: an N-dimensional probability density function transform for climate model simulations of multiple variables. *Climate Dynamics* 50 doi: [10.1007/s00382-017-3580-6](https://doi.org/10.1007/s00382-017-3580-6).
- Chandler RE** (2020) Multisite, multivariate weather generation based on generalised linear models. *Environmental Modelling & Software* 134, 104867. issn: 1364–8152. doi: [10.1016/j.envsoft.2020.104867](https://doi.org/10.1016/j.envsoft.2020.104867). <https://www.sciencedirect.com/science/article/pii/S1364815220309245>.
- Christensen JH**, et al. (2008) On the need for bias correction of regional climate change projections of temperature and precipitation. *Geophysical Research Letters* 35. <https://api.semanticscholar.org/CorpusID:128752652>.
- Christensen NS**, et al. (2004) The effects of climate change on the hydrology and water resources of the Colorado River basin. *Climatic Change* 62, 337–363. <https://api.semanticscholar.org/CorpusID:53533021>.
- Ciais P**, et al. (2005) Europe-wide reduction in primary productivity caused by the heat and drought in 2003. *Nature* 437, 529–33. doi: [10.1038/nature03972](https://doi.org/10.1038/nature03972).
- Defrance D** et al. (2017) Consequences of rapid ice sheet melting on the Sahelian population vulnerability. *Proceedings of the National Academy of Sciences* 114(25), 6533–6538. doi: [10.1073/pnas.1619358114](https://doi.org/10.1073/pnas.1619358114). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1619358114>. <https://www.pnas.org/doi/abs/10.1073/pnas.1619358114>.
- Dekens L** et al. (2017) Multivariate distribution correction of climate model outputs: A generalization of quantile mapping approaches. *Environmetrics* 28. <https://api.semanticscholar.org/CorpusID:126106313>.
- Déqué M** (2007) Frequency of precipitation and temperature extremes over France in an anthropogenic scenario: Model results and statistical correction according to observed values. *Global and Planetary Change* 57, 16–26. <https://api.semanticscholar.org/CorpusID:128409261>.
- Eden JM**, et al. (2012) Skill, correction, and downscaling of GCM-simulated precipitation. *Journal of Climate* 25, 3970–3984. <https://api.semanticscholar.org/CorpusID:128629641>.
- Eyring V**, et al. (2016) Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development* 9(5), 1937–1958. doi: [10.5194/gmd-9-1937-2016](https://doi.org/10.5194/gmd-9-1937-2016). <https://gmd.copernicus.org/articles/9/1937/2016/>.
- François, B**, et al. (2020) Multivariate bias corrections of climate simulations: which benefits for which losses? *Earth System Dynamics* 11, 537–562. doi: [10.5194/esd-11-537-2020](https://doi.org/10.5194/esd-11-537-2020). <https://esd.copernicus.org/articles/11/537/2020/>.
- Gleick PH** (1989) Climate change, hydrology, and water resources. *Reviews of Geophysics* 27, 329–344. <https://api.semanticscholar.org/CorpusID:128418811>.
- Grouillet B**, et al. (2015) Sensitivity analysis of runoff modeling to statistical downscaling models in the western Mediterranean. *Hydrology and Earth System Sciences* 20, 1031–1047. <https://api.semanticscholar.org/CorpusID:55302374>.
- Hagemann S**, et al. (2011), Impact of a statistical bias correction on the projected hydrological changes obtained from three GCMs and two hydrology models. *Journal of Hydrometeorology* 12, 556–578.
- Hawkins E**, et al. (2013) Calibration and bias correction of climate projections for crop modelling: An idealised case study over Europe. *Agricultural and Forest Meteorology* 170, 19–31.
- Hersbach H**, et al. (2020) The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* doi: [10.1002/qj.3803](https://doi.org/10.1002/qj.3803).
- Ho CK**, et al. (2012) Calibration strategies a source of additional uncertainty in climate change projections. *Bulletin of the American Meteorological Society* issn: 00030007. doi: [10.1175/2011BAMS3110.1](https://doi.org/10.1175/2011BAMS3110.1).
- Hourdin F**, et al. (2020) LMDZ6A: The Atmospheric Component of the IPSL Climate Model With Improved and Better Tuned Physics". In: *Journal of Advances in Modeling Earth Systems* 12(7), e2019MS001892 [10.1029/2019MS001892](https://doi.org/10.1029/2019MS001892), <https://doi.org/10.1029/2019MS001892>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS001892>. url: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001892>.
- Hüllermeier E and Waegeman W** (2019) Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning* 110, 457–506. <https://api.semanticscholar.org/CorpusID:216465307>.
- Iizumi T, Yokozawa M and Nishimori M** (2011) Probabilistic evaluation of climate change impacts on paddy rice productivity in Japan. *Climatic Change* 107, 391–415.
- Kallache M**, et al. (2010) Nonstationary probabilistic downscaling of extreme precipitation. *Journal of Geophysical Research* 116 <https://api.semanticscholar.org/CorpusID:67764166>.
- Kalnay E**, et al. (1996) The NCEP/NCAR 40-Year reanalysis project. *Renewable Energy*. <https://api.semanticscholar.org/CorpusID:124135431>.
- Li H, Sheffield J and Wood E**. Bias correction of monthly precipitation and temperature fields from Intergovernmental Panel on Climate Change AR4 models using equidistant quantile matching. *Journal of Geophysical Research* 115 (2010).
- Liu PJ**, et al. Generating wikipedia by summarizing long sequences. *ArXiv abs/1801.10198* (2018). <https://api.semanticscholar.org/CorpusID:3608234>.
- Lurton T**, et al. (2020) Implementation of the CMIP6 Forcing Data in the IPSL-CM6A-LR Model. *Journal of Advances in Modeling Earth Systems* 12(4), e2019MS001940 [10.1029/2019MS001940](https://doi.org/10.1029/2019MS001940), doi: <https://doi.org/10.1029/2019MS001940>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019MS001940>. url: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001940>.

- Maraun D** (2013) Bias correction, quantile mapping, and downscaling: revisiting the inflation issue. *Journal of Climate* 26, 2137–2143.
- Maraun D and Widmann M** (2018) *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press, doi: [10.1017/9781107588783](https://doi.org/10.1017/9781107588783).
- Maraun D, et al.** (2015) VALUE: a framework to validate downscaling approaches for climate change studies. *Earth's Future* 3(1) 1–14. doi: <https://doi.org/10.1002/2014EF000259>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2014EF000259>. url: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014EF000259>.
- Marx W, Haunschild R and Bornmann L** (2021) Heat waves: a hot topic in climate change research. *Theoretical and Applied Climatology* 146, 781–800. <https://api.semanticscholar.org/CorpusID:235652213>.
- Mehrotra R and Sharma A** (2019) A resampling approach for correcting systematic spatiotemporal biases for multiple variables in a changing climate. *Water Resources Research* 55, 754–770. <https://api.semanticscholar.org/CorpusID:134017482>.
- Nivron O, Parthipan R and Wischik D** (2023) Taylorformer: probabilistic predictions for time series and other processes. *ArXiv abs/2305.19141*. <https://api.semanticscholar.org/CorpusID:258967520>.
- OECD** (2015) *The Economic Consequences of Climate Change*. 140. doi: <https://doi.org/10.1787/9789264235410-en>. <https://www.oecd-ilibrary.org/content/publication/9789264235410-en>.
- Palmer TN and Weisheimer A.** (2011) Diagnosing the causes of bias in climate models—why is it so hard? *Geophysical & Astrophysical Fluid Dynamics* 105(2–3), 351–365. doi: [10.1080/03091929.2010.547194](https://doi.org/10.1080/03091929.2010.547194). eprint: <https://doi.org/10.1080/03091929.2010.547194>.
- Panofsky HA and Brier GW** (1968) *Some Applications of Statistics to Meteorology*. Earth and Mineral Sciences Continuing Education, College of Earth and Mineral Sciences.
- Piani C and Haerter JO** (2012) Two dimensional bias correction of temperature and precipitation copulas in climate models. *Geophysical Research Letters* 39(20) doi: <https://doi.org/10.1029/2012GL053839>. eprint: <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2012GL053839>. <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2012GL053839>.
- Piao S, et al.** (2010) The impacts of climate change on water resources and agriculture in China. *Nature* 467, 43–51. url: <https://api.semanticscholar.org/CorpusID:4410647>.
- O'Reilly CH, Befort DJ and Weisheimer A** (2020) Calibrating large-ensemble European climate projections using observational data. *Earth System Dynamics* 11(4) 1033–1049. doi: [10.5194/esd-11-1033-2020](https://doi.org/10.5194/esd-11-1033-2020). <https://esd.copernicus.org/articles/11/1033/2020/>.
- Richard SJT** (2021) The economic impact of weather and climate. *Climate Action eJournal* <https://api.semanticscholar.org/CorpusID:228207225>.
- Richard SJT** (2018) The economic impacts of climate change. *Review of Environmental Economics and Policy* 12, 4–25. <https://api.semanticscholar.org/CorpusID:159009492>.
- Robin Y, et al.** (2019) Multivariate stochastic bias corrections with optimal transport. *Hydrology and Earth System Sciences* 23(2), 773–786. doi: [10.5194/hess-23-773-2019](https://doi.org/10.5194/hess-23-773-2019). url: <https://hess.copernicus.org/articles/23/773/2019/>.
- Robin Y, et al.** (2019) Multivariate stochastic bias corrections with optimal transport. *Hydrology and Earth System Sciences* 23(2), 773–786. doi: [10.5194/hess-23-773-2019](https://doi.org/10.5194/hess-23-773-2019). url: <https://hess.copernicus.org/articles/23/773/2019/>.
- Robin Y and Vrac M** (2021) Is time a variable like the others in multivariate statistical downscaling and bias correction? *Earth System Dynamics*. <https://api.semanticscholar.org/CorpusID:236726767>.
- Robin Y and Vrac M** (2023) *SBCK: Statistical Bias Correction Kit*. Version 1.0.0. R package version 1.0.0, Package source: [SBCK_1.0.0.tar.gz](https://github.com/yrobink/SBCK). R Foundation. CRAN: Comprehensive R Archive Network, Sept. <https://github.com/yrobink/SBCK>.
- Rumelhart DE, Hinton GE and Williams RJ** (1986) Learning representations by back-propagating errors. *Nature* 323(6088), 533–536.
- Salvi K, et al.** (2011) Statistical downscaling and bias correction for projections of Indian rainfall and temperature in climate change studies. *International Conference on Environmental and Computer Science, IPCBEE vol.19*. IACSIT Press, Singapore.
- Schmidli J, Frei C and Vidale PL** (2006) Downscaling from GCM precipitation: a benchmark for dynamical and statistical downscaling methods. *International Journal of Climatology* 26. <https://api.semanticscholar.org/CorpusID:128686717>.
- Stainforth D, et al.** (2007) Confidence, uncertainty and decision-support relevance in climate predictions. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365, 2145–2161.
- Summers T, et al.** (2022) Localized impacts and economic implications from high temperature disruption days under climate change. *Climate Resilience and Sustainability* 1(2)e35. doi: <https://doi.org/10.1002/cli2.35>. eprint: <https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/cli2.35>. url: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/cli2.35>.
- Teutschbein C, et al.** (2012) Bias correction of regional climate model simulations for hydrological climate-change impact studies: Review and evaluation of different methods. *Journal of Hydrology* 456, 12–29. url: <https://api.semanticscholar.org/CorpusID:129017016>.
- Themesl M, Gobiet A and Leuprecht A** (2011) Empirical-statistical downscaling and error correction of daily precipitation from regional climate models. *International Journal of Climatology* 31, 1530–1544.
- Turco M, et al.** (2017) Bias correction and downscaling of future RCM precipitation projections using a MOS-Analog technique. *Journal of Geophysical Research* 122, 2631–2648.
- Vaswani A, et al.** (2017) Attention is all you need. *Advances in Neural Information Processing Systems*. <https://api.semanticscholar.org/CorpusID:13756489>.

- 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. doi: 10.5194/hess-22-3175-2018. <https://hess.copernicus.org/articles/22/3175/2018/>.
- Vrac M and Thao S.** (2020) R2D2 v2.0: accounting for temporal dependences in multivariate bias correction via analogue rank resampling. *Geoscientific Model Development* 13(11), 5367–5387. doi: 10.5194/gmd-13-5367-2020. <https://gmd.copernicus.org/articles/13/5367/2020/>.
- Vrac M, et al.** (2012) Dynamical and statistical downscaling of the French Mediterranean climate: uncertainty assessment. *Natural Hazards and Earth System Sciences* 12. (9), pp. 2769–2784. doi: 10.5194/nhess-12-2769-2012. <https://nhess.copernicus.org/articles/12/2769/2012/>.
- Vrac M and Friederichs** (2015) Multivariate—intervariable, spatial, and temporal—bias correction*. *Journal of Climate* 28 218–237. <https://api.semanticscholar.org/CorpusID:128702918>.
- Williams RJ and Zipser D.** (1989) A learning algorithm for continually running fully recurrent neural networks. *Neural Computation* 1(2), 270–280.
- Wood A, et al.** (2002) Long-range experimental hydrologic forecasting for the eastern United States. *Journal of Geophysical Research* 107, 4429.
- Xu C-y** (1999) From GCMs to river flow: a review of downscaling methods and hydrologic modelling approaches. *Progress in Physical Geography: Earth and Environment* 23.(2), 229–249. doi: 10.1177/030913339902300204. eprint: <https://doi.org/10.1177/030913339902300204>. url: <https://doi.org/10.1177/030913339902300204>.

Appendix

A. Distribution shift

In the context of climate change, we are often interested to see if a specific location and a specific physical variable (e.g., daily maximum temperatures) have seen a change in its distribution.

Here, we show that for selected months, the two chosen locations, Tokyo and Abuja, had experienced a shift in their distributions. It is apparent both in the observed record (ERA5) and in the climate model (IPSL), and their respective changes seem to be qualitatively different.

QQ Plot: observed daily max temperature (Tmax) > 90th quantile in Tokyo (1968–1988) and (1989–2008)




Figure 6. Changes in the distribution of daily maximum temperatures above the 90th quantile in Tokyo: This plot compares the distribution of temperatures between two periods: 1968–1988 (X-axis) and 1989–2008 (Y-axis).

QQ Plot: observed daily max temperature (Tmax) > 90th quantile in Abuja (1968-1988) and (1989-2008)




Figure 7. Changes in the distribution of daily maximum temperatures above the 90th quantile in Abuja: This plot compares the distribution of temperatures between two periods: 1968–1988 (X-axis) and 1989–2008 (Y-axis).




Figure 8. The distribution of daily observed (ERA5) maximum temperatures (°C) in Tokyo, Japan had changed from 1968–1988 (blue histogram) to 1989–2009 (orange histogram) in February (left) and July (right).




Figure 9. The distribution of daily observed (ERA5) maximum temperatures ($^{\circ}\text{C}$) in Abuja, Nigeria had changed from 1968–1988 (blue histogram) to 1989–2009 (orange histogram) in February (left) and July (right).




Figure 10. The distribution of modeled (IPSL) ensemble average of maximum daily temperatures in Tokyo, Japan had changed between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in February (left) and July (right). Qualitatively, the change is different than that in the observed record (see Figure 8).




Figure 11. The distribution of modeled (IPSL) ensemble average of maximum daily temperatures in Abuja, Nigeria had changed between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in February (left) and July (right). Qualitatively, the change is different than that in the observed record (see Figure 9).




Figure 12. The distribution of observed (ERA5) daily maximum temperature above 22 °C in Tokyo had shifted between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in July (left) and August (right).




Figure 13. The distribution of observed (ERA5) daily maximum temperature above 30 °C in Abuja had shifted between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in January (left) and February (right).




Figure 14. The distribution of modeled (IPSL) daily maximum temperature above 22 °C in Tokyo had shifted between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in January (left) and February (right). Qualitatively, the change is different than that in the observed record (see Figure 12).




Figure 15. The distribution of modeled (IPSL) daily maximum temperature above 30 °C in Abuja had shifted between 1968–1988 (blue histogram) and 1989–2009 (orange histogram) in January (left) and February (right). Qualitatively, the change is different than that in the observed record (see Figure 13).

B. Motivation for the Taylorformer temporal BC

Here, we present the motivation to reformulate BC as a time-indexed regression task. Further, we discuss why we thought it could address the time asynchronicity between climate models and observations and how it can resolve local-scale variability.

First, we show that some standard BC methods, such as mean shift, can already be written as a time-indexed regression model. However, since their outputs cannot be taken as literal values at a specific time point, we progress to more flexible models that can.

B.1. A univariate BC as a time-indexed regression model

In our method, we use a regression model to perform the BC. However, other BC methods can also be performed by regression. In the following, we examine the mean shift (Xu, 1999) (used, for example, by Mora et al. (2017) for heatwave mortality impacts).

In a mean shift, observations and climate model outputs are collected over a reference period, and the difference between their means is calculated. The difference is then added to future climate projections (hereafter, we call this execution the “procedural approach”).

Alternatively, let O_i be the random variable representing the temperature readouts at time-point t_i for observations and g_i the corresponding climate model value. Now, the following probability model can be written:

$$O_i = g_i + c + E_i \tag{4}$$

where $E_i \sim \text{Normal}(0, \sigma^2)$ and E_i is independent of E_j for $i \neq j$ and c is a learned parameter. To find c , we can use standard linear regression packages. Standard linear regression uses maximum likelihood as its objective to learn the parameters. We show in Appendix 8 the equivalence of the procedural approach and the probability model. As a trivial extension, the mean and variance shift probability model is

$$O_i = b \times g_i + c + E_i \tag{5}$$

where both b, c are learned parameters using standard linear regression (again using the maximum likelihood objective).

This setup of the two equations above is unusual in the BC literature since the index i is the same on both sides of the equation. Generally, BC practitioners avoid such a formulation since we know that climate models and observations are not in temporal synchrony (Maraun & Widmann, 2018, p. 137). Nevertheless, we can write it in this way and arrive at the same results as the “procedural approach.”

This view shows that we can use regression models for BC (and, in a sense, we already do implicitly). For the BC practitioner, it may be of interest that this view also provides a way to treat mean shift and mean–variance shift as stochastic BC methods (here, the stochasticity comes from what is known in statistics and ML literature as aleatoric uncertainty (Hüllermeier & Waegeman, 2019).

After performing the regression, we will get, for each random variable, O_i , a distribution with a mean value equaling $g_i + c$. In this regression model, the mean value of O_i at time t_i should not be taken to be the literal value at that exact time point due to temporal asynchrony. In other words, if tested using MSE on an unseen hold-out data, the value of $\frac{1}{M} \sum_{i=1}^M (o_i - [g_i + c])^2$ will be large, i.e., the model performs poorly.

The fact that the distribution obtained by O_i at time t_i should not be taken to be the literal values at that exact time point leads us to search for more flexible regression models where it may be.

Even if the two time series were synchronized, it is clear that we cannot use it to correct time-dependent statistics since the random variables O_i and O_j are independent of each other for any $i \neq j$. This is equivalent to throwing away all temporal correlation information and relying solely on the temporal correlation of the climate model.

B.2. Outputs with a specific time index can be meaningful

In weather forecasting, when we say output with a specific time-index is meaningful, we mean that the estimated value \hat{o}_i of the random variable O_i is very close (by MSE) to the observed value o_i and its confidence interval is very narrow. To be “meaningful,” however, in a time-series context on a longer time horizon, we do not need to have a narrow confidence band or a small MSE value. On a climatic time horizon, we cannot precisely specify the value obtained in Tokyo, Japan, on January 1st, 2050. However, we can still look at the distribution at that specific time point. One added benefit to this approach is that we can account for temporal correlations, i.e., what happened on day t can affect the possible values at day $t + 1$.

The fact that some probability models can be useful for long-time horizons, even when looking at a specific time index, is not specific to BC. As an illustrative example, we will discuss a general prediction task for time series using a Gaussian Process (GP) regression model. The GP is specified by what is known as a Radial Basis Function (RBF) kernel. In Figure 16, we show that when extrapolating further from the training data points (“+”), our mean prediction (solid blue line) converges to zero. Let us look at the marginal distribution of the target variable Y_{80} with $t_{80} = 1.2$ (roughly, data between the dashed green lines).




Figure 16. The task involves predicting future values using training points (“+”) located left of the vertical black line. The predicted mean (solid blue line) aligns closely with these points and quickly reverts to zero beyond this boundary. The light blue shading indicates the 95% confidence interval. Sample trajectories from the predicted probability model are shown as colored lines to the right of the black line. These demonstrate that target variables distant from the training data maintain the same marginal distribution, mean, and variance, as evidenced by the vertical slice marked by green dashed lines. This time-series view highlights correlations among points, offering insights not visible from a simple histogram of training points.




Figure 17. Predictions are made using training points (“+”) left of the vertical black line. The predicted mean (solid blue line) aligns accurately with these points and, upon extrapolation to the right, adapts to typical “seasonality” behaviors instead of becoming non-informative (refer to Figure 16). The light blue shading indicates the 95% confidence interval for each x -axis value. Various sample trajectories from the probability model are shown as colored lines, while the dashed green lines illustrate the marginal distribution at corresponding x -axis values.

We will get $Y_{80} \sim \text{Normal}(0, 1.52)$, and $Y_{90} \sim \text{Normal}(0, 1.52)$ as well. There are two takeaways from this model: (1) random variables further away from the data will revert to the typical behavior (mean zero), and their confidence interval will widen in comparison to regions close to the training data. This is seen clearly in the region to the left of the vertical black line (the X -axis value is zero) and to its right. However, they still provide useful information, namely the typical behavior, when viewed with a specific time index. And (2) while the marginal distribution is the same correlation information shown by the wiggly lines to the right of the vertical black line.

As a second example, consider another GP shown in Figure 17: the random variables’ mean (solid blue line) far away from the training data (“+”) again reverts to the typical behaviour. However, in this case, the typical behavior includes the “seasonality” seen in the past. (See Appendix 13 for the exact details of the two GP models above.)

These are two pictures that serve as a motivation for the temporal BC task. Similarly to the pictures of the GP, we would expect the marginal distribution of random variables far away from the training data to describe a typical behavior such as seasonal, yearly, or decadal means. But, an additional benefit would be learning the correlation structure between variables. To learn the correlation structure, we will need a regression model that allows for dependence between points (as opposed to a mean shift or an EQM model).

B.3. Time indexed BC can correct systematic time asynchronicities

First, let us look at a simple toy example of time asynchronicity. In the left panel of Figure 18, we show two identical time series in which we added a mean bias to the pseudo observations (blue points) and then horizontally shifted the whole time series to the right. Here, the horizontal shift is the “time asynchronicity.” The red points represent the pseudo-climate model data. If we were to use 1D histograms over the same reference period, as is typical in univariate BC, we would not be able to infer that the two time series are identical. Differently, a time-indexed probability model like a GP can learn the time-shift from data, as shown in the right panel of Figure 18. Note that the input to the GP includes both (past and future) pseudo-climate model and past observations to make the correction possible.

As a second example, imagine the number of heatwaves is consistently two times longer in the climate model than in the observational record. We can then manually construct a probability model. Our probability model copies the climate model time series; if the time series passes above 30°C for multiple days in a row (“heatwave”), cut the “heatwave” time series in half and add a bias a and some noise, E_i , to get the corresponding estimated observation; otherwise, assign 0.




Figure 18. This figure uses a toy example of a Gaussian Process-based BC regression (details in Appendix 13) to illustrate correction of systematic time asynchronicities between “climate model” and observations. Left panel: The pseudo observations (blue points) are both vertically and horizontally shifted relative to the pseudo climate model (red points). Right panel: The GP model’s mean (golden solid line) successfully captures the vertical and horizontal shifts, as shown by the alignment of the pseudo observations (solid blue line) with the pseudo climate model’s full time-series (solid red line). Traditional quantile mapping approaches with a fixed reference period would fail to recognize that the processes are identical, barring a mean shift.

The mathematical notation for the associated probability model is not important here, but the main takeaway is that the value at time point t_i , o_i , is some known deterministic function f of the (past, future) values from the climate model. Since we added bias a and a noise term, we can write the model mathematically as

$$O_k | inp_k \sim \text{Normal}(f(inp_k) + a, \sigma^2) \tag{6}$$

where $inp_k = (g_1, \dots, g_M, k)$. That is, again, an easy probability model with a learnable parameter a which we can learn using standard linear regression libraries.

If in addition to the number of heatwaves setup, we have a consistent mean shift of one degree (°C) between the climate model and observations, we can learn this shift by adding to the input, inp_k , past observations, o_1, \dots, o_n , where $n < K$. From this example, we see that the values from the climate model in the (past, future) and past observations will be useful in the probability model we are constructing.

B.4. Time-indexed BC coupled with downscaling can resolve local scale variability

If the resolution of the climate model is coarser than that of the observed dataset, one grid point in the climate model is the average behavior of multiple grid points in the observed data. If we use EQM or mean shift, for example, high values in the projections of a GCM will map to high estimated observed values, regardless of the specific behavior at the corresponding grid point in the observed data. For quantile-mapping approaches, this is also known as the inflation issue (Maraun, 2013) (See more details in Maraun & Widmann, 2018, pp. 189–192)

To have a model that can differentiate between the local scale grid points, we can simply add the time series of the observed values as input. This is another reason for adding past observations to our regression model.

B.5. Tying it all together

In the context of BC, we can perform a regression model. Due to a lack of temporal synchrony, a regression model (like mean shift) with one-to-one correspondence would perform poorly on a hold-out set, and its values should not be taken literally. However, a

regression model with a specific time index that does not assume independence between target variables, like a GP (or other classic AR model), can a) describe the typical distributional behavior of points far away from the training data without throwing away temporal information, and b) since it receives as input many values from different time points (past, future) climate model and past observations can find patterns that are not a one-to-one correspondence and potentially fix time asynchronicities (such as in Figure 18). Furthermore, suppose the climate model is coarser in resolution than the observed data (one climate model grid point to many observed grid points). In that case, the regression model can potentially differentiate between the grid points and adapt to local variability if its input includes past observations.

We can write the GP version of the temporal BC probability model as

$$O_k | inp_k \sim \text{Normal}(f(inp_k), g(inp_k)) \tag{7}$$

and M being the size of the sequence. Furthermore, we denote with t' the time points of the climate model, that are potentially different than the observed time-points. Additionally, here $inp_k = ((g_1, t'_1), \dots, (g_M, t'_M), (o_1, t_1), \dots, (o_{k-1}, t_{k-1}))$, f and g are functions of the data defined by the GP.

If the systematic patterns of time asynchronous were clear to the human eye or by standard analysis, we could construct a GP model with the appropriate equations. However, since for the BC task, we do not know how to specify these interactions, we turn to a more flexible neural network model that can learn these interactions from the data.

C. Linear regression mean shift derivation

Estimation of the parameters in linear regression is done using maximum likelihood as the objective. Using maximum likelihood and denoting \mathbf{g} to refer to the set of all climate model outputs, we have

$$\begin{aligned} \max_{\theta} P(O_1, \dots, O_N | \mathbf{g}) &\stackrel{\text{independence}}{=} \prod_{i=1}^N P(O_i | g_i) \\ &\stackrel{\text{normal i.v.}}{=} \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(o_i - [g_i - c])^2}{2\sigma^2}\right) \\ &\xrightarrow{\log} \sum_{i=1}^N \log(1) - \log \sqrt{2\pi\sigma^2} - \frac{(o_i - [g_i - c])^2}{2\sigma^2} \tag{8} \\ &\frac{\partial}{\partial c} \xrightarrow{\text{derivative}=0} - \sum_{i=1}^N ((o_i - [g_i - c])) = 0 \\ &= \frac{1}{N} \sum_{i=1}^N o_i - \frac{1}{N} \sum_{i=1}^N g_i \end{aligned}$$

D. From Transformer decoder to Taylorformer

D.1. Transformer decoder for language modeling

One component in recent successes of large language models, like GPT (Brown et al., 2020), is their Transformer decoder architecture. We will now give an abstract description of the architecture—the specifics can be found in (Liu et al., 2018), but the abstract view, we believe, makes it easier to explain the salient parts of the model and makes it clear which parts can be replaced with ease.

For language modeling, our dataset comprises N sequences, each comprised of (word, position) pairs (W_i, P_i) for $i \in 1 \dots K_s$. K_s is the length of sequence s . Both W_i and P_i are scalars. We will take D to denote the total number of unique words in our dataset.

The goal of language modeling is to classify the next word in a sentence based on previous words.

D.1.1. Feature engineering

In the raw data, P_i is a scalar. One of the ingenious of the Transformer decoder is the extraction of features from P_i ; these are prepared before training and are inspired by Fourier analysis. The formula is

$$P'(i, 2l) = \sin\left(\frac{P_i}{(10000)^{2l/d}}\right) \tag{9}$$

$$P'(i, 2l+1) = \cos\left(\frac{P_i}{(10000)^{2l/d}}\right) \tag{10}$$

l is the dimension and after the features are extracted $P'_i \in \mathbf{R}^d$.

D.1.2. One attention layer

As a first step, we map each word W_i to a vector of learnable weights $W'_i \in \mathbf{R}^h$. $W'_i = W'_j$ if and only if $W_i = W_j$. Secondly, we concatenate W'_i and P'_i and pass them through three separate non-linear functions (could be multiple layers of a neural network) f_θ, g_θ and h_θ , let us call the resulting vectors Q_i, K_i, V_i all of which are in \mathbf{R}^d .

Attention weights. Are then simply constructed as $A_{i,j} = \text{Softmax}_{j \in 1, \dots, i} \left(\sum_{\ell=1}^d Q_{i,\ell} K_{\ell,j} \right)$.

Prediction. for the probabilities for all classes (all possible words in the dataset) for target variable W_i after one layer can then be made by $\pi_i(\theta) = \text{Softmax}_{d \in 1, \dots, D} B_\theta \left(\sum_{j=1}^{i-1} A_{i-1,j} \times V_j \right)$. $B_\theta : \mathbf{R}^d \rightarrow \mathbf{R}^D$ is a nonlinear function with dimension output equal to the number of unique words in the dataset D , and $\pi_i(\theta) \in \mathbf{R}^D$. Here, we assumed the target word is distributed according to a Categorical random variable with D categories. We can write it as $W_i \sim \text{Categorical}(\pi_{i,1}(\theta), \dots, \pi_{i,D}(\theta))$.

D.1.3. Multiple attention layers

Each attention layer after the first one assumes the same structure: take the output (before the last Softmax operator) from the previous layer to be the new values for vectors Q_i, K_i, V_i and repeat the calculation of attention weights and prediction (excluding the last Softmax). In the last layer, add the Softmax to get a probability for a word at position i .

D.1.4. Multiple attention heads

It is common to split the procedure within each attention layer into multiple heads (or branches). This works by taking the vectors Q_i, K_i, V_i and splitting them into H heads such that V_i is equal to the concatenation of V_i^h for $h \in 1, \dots, H$ and $V_i^h \in \mathbf{R}^{d/H}$ and a similar split is performed for K_i and Q_i . Then, we can follow the same procedure as before for attention weights for each head h to get $A_{i,j}^h$. After we calculate $\sum_{j=1}^{i-1} A_{i-1,j}^h \times V_j^h$, we can concatenate all the heads and continue as before.

D.1.5. Training objective

The training objective will be to maximize the log-likelihood of the data. For one sentence with words W_1, \dots, W_{K_s} , this means

$$\text{maximise}_\theta \text{Log} P(W_1, \dots, W_{K_s}, \theta) \underbrace{=}_{\text{chain rule}} \text{maximise}_\theta \text{Log} (P(W_{K_s} | W_{1:K_s-1}, \theta) \dots P(W_2 | W_1, \theta) P(W_1 | \theta)),$$

where we have used the chain rule to decompose the joint probability and θ refers to the parameters we want to learn via gradient descent using our model. Now, we just need to replace each term in the equation with the categorical distribution. If the true word, W_i at position i is $w_i \in \mathbf{N}$, we will denote its predicted probability as π_{i,w_i} , then we can write the log-likelihood maximization objective as

$$\text{maximise}_\theta \text{Log} \pi_{1,w_1}(\theta) \times \dots \times \pi_{K_s,w_{K_s}}(\theta) \tag{11}$$

D.1.6. Inference

At the inference stage, we can simply fix $W_1 = 0$ where the 0-th index indicates the first word being the token < sos > (start-of-sentence), we then do a forward pass of our model to get π_2 , we then sample $(W_2 | W_1 = 0) \sim \text{Categorical}(\pi_{2,0}(\theta), \dots, \pi_{2,D}(\theta))$. Next, we make another forward pass now with the fixed value of W_1 and the sampled value of $(W_2 | W_1 = w_1)$ to get $\pi_3(\theta)$. We continue this procedure until the output of the sampling gives us the token < eos > (end-of-sentence).

D.1.7. Language sequences versus continuous processes

Language sequences follow a discrete regular grid—at each position p we are interested in the word at position $p + 1$ as opposed to predicting k steps ahead, i.e., what will be the word at position $p + k$ given the words up to position p . For this reason, in the equations above, when predicting the next word, we do not require information about the position to be predicted next. This situation may fit some continuous processes, but there are other situations for continuous processes, like interpolation, long-term prediction and missing values, that would not work well with the same setup. We first introduce the Transformer decoder for regular grid continuous processes and then extend it to the more general case.

D.2. Transformer decoder for regular-grid continuous processes

A simple adaptation is needed for translating the Transformer decoder from a language model to a continuous processes model. Briefly, we just need to adapt the feature engineering, remove the Softmax operator from the prediction stage and choose a different distribution for our targets. We do not need to change anything about the “Multiple attention layers” and “Multiple heads” sections.

D.2.1. Feature engineering

In the raw data, now we have time point, $t_i \in \mathbf{R}$, instead of a position, $P_i \in \mathbf{N}$. The adapted formula we suggest is

$$P'(t_i, 2l) = \sin \left(\frac{t_i / \delta_t}{(t_{\max} / \delta_t)^{2l/d}} \right) \tag{12}$$

$$P'(t_i, 2i + 1) = \cos\left(\frac{t_i/\delta_t}{(t_{\max}/\delta_t)^{2t_i/d}}\right) \tag{13}$$

So after the features are extracted $P'_i \in \mathbf{R}^d$. Where t_{\max}, δ_t are hyperparameters.

D.2.2. One attention layer

In the continuous case, we have real-valued random variables O_i with their observed values being o_i at time-point t_i . In language modeling, we mapped the index of word W_i to a vector of learnable weights. Here, we can just pass each value o_i such that $O'_i \in \mathbf{R}^h$, where $O'_i = A_\theta(o_i)$ where A_θ denotes nonlinear neural network layers. Then, we proceed as before by concatenating O'_i and P'_i and pass them through three separate nonlinear functions (could be multiple layers of a neural network) f_θ, g_θ , and h_θ , let us call the resulting vectors Q_i, K_i, V_i all of which are in \mathbf{R}^d .

Attention weights. Are calculated in exactly the same manner as before, i.e., $A_{i,j} = \text{Softmax}_{j \in 1, \dots, i} \left(\sum_{\ell=1}^d Q_{i,\ell} K_{\ell,j} \right)$.

Prediction. for the continuous case differs only in the last layer; instead of a last layer of *Softmax* to get class probabilities (used in the language model), we can use for the last layer a linear layer $B_\theta \mathbf{R}^d \rightarrow \mathbf{R}^2$, such that $\mu_i(\theta), \sigma_i(\theta) = B_\theta \left(\sum_{j=1}^{i-1} A_{i-1,j} \times V_j \right)$. Here, we assumed the target random variable is distributed according to a Normal random variable with two hyperparameters μ, σ . We can write it as $O_i \sim \text{Normal}(\mu_i(\theta), \sigma_i(\theta))$.

D.2.3. Training objective

The training objective will be to maximize the log-likelihood of the data. For one sequence with observations O_1, \dots, O_{K_s} , this means

$$\text{maximise}_\theta \text{Log} P(O_1, \dots, O_{K_s}, \theta) \underset{\text{chain rule}}{=} \text{maximise}_\theta \text{Log}(p(O_{K_s} | O_{1:K_s-1}, \theta) \cdots p(O_2 | O_1, \theta) p(O_1 | \theta)),$$

where we have used the chain rule to decompose the joint probability and θ refers to the parameters we want to learn via gradient descent using our model. Now, we just need to replace each term in the equation by the Normal distribution. If the true value, O_i at position i is $o_i \in \mathbf{R}$, we will denote its estimated Normal probability hyperparameters as $\mu_i(\theta), \sigma_i(\theta)$, then we can write the log-likelihood maximization objective as

$$\text{maximise}_\theta \text{Log} \sum_{i=1}^{K_s} \frac{1}{\sqrt{2\pi\sigma_i(\theta)}} \exp - \frac{(o_i - \mu_i(\theta))^2}{2\sigma_i(\theta)^2} \tag{14}$$

This could be further simplified to give

$$\text{maximise}_\theta \sum_{i=1}^{K_s} \text{Log} \left(\frac{1}{\sqrt{2\pi\sigma_i(\theta)}} \right) - \frac{(o_i - \mu_i(\theta))^2}{2\sigma_i(\theta)^2} \text{ where yellow box indicates the squared error} \tag{15}$$

If we choose our model to output just the mean in the prediction layer, the training objective becomes the MSE. Note that we could choose to provide our model with initial “information” (content) by fixing few of the random variables to their observed values and then maximizing the likelihood conditioned on the content—this is what we do in Bias-Correction.

D.2.4. Inference

At the inference stage, we can simply fix O_1 to the last observed value during training (this is just one option out of many); we then do a forward pass of our model to get $(\mu_2(\theta), \sigma_2(\theta))$; and we then sample $(O_2 | O_1 = o_1) \sim \text{Normal}(\mu_2(\theta), \sigma_2(\theta))$. Next, we make another forward pass now with the fixed value of O_1, o_1 , and the sampled value of $(O_2 | O_1 = o_1)$ to get $(\mu_3(\theta), \sigma_3(\theta))$. We continue this procedure until desired.

D.3. Transformer decoder for irregular-grid continuous processes

In a Transformer decoder for a regular grid, we described the inner part of the prediction (hereafter “weighted attention”) equation as $B_\theta \left(\sum_{j=1}^{i-1} A_{i-1,j} \times V_j \right)$. Notice that this is a prediction for the i -th timestamp, but none of the elements involved in the equation has the data for the next timestamp.

To adapt it to the continuous case, we can a) add different inputs to the Q_i vector and b) use “weighted attention.”

Adapted Q_i . The input to Q_i will be the same as before for values at time points t_i we want to provide the model in advance and a concatenation of P'_i and 0 for values at other time points t_i we wish to predict, i.e., we mask the associated value of O'_i .

Adapted “weighted attention”. The equation of the “weighted attention” for O_i will now be $B_\theta \left(\sum_{j=1}^{i-1} A_{i,j} \times V_j \right)$. Notice that we have changed $A_{i-1,j}$, which did not contain information about the next time stamp, to $A_{i,j}$, which contains the information about t_i through multiple nonlinear layers leading to the vector Q_i . The sum still runs up to $i - 1$ since K_i and V_j stayed the same and did contain information about the actual value of O_i , which we do not want to leak. Since Q_i contains information about the next

desired time point, we can ask for a long-term prediction in the future or in the past, and we do not need the time points to be equally spaced.

D.4. Taylorformer decoder

In the Taylorformer, we have three key differences to the continuous Transformer decoder: (1) a new prediction layer, (2) different inputs to the vectors Q_i, K_i, V_i , and (3) an additional attention mechanism.

New prediction layer: Instead of predicting the hyper parameters ($\mu_i(\theta), \sigma_i(\theta)$) for O_i , we predict the parameters ($\mu_i(\theta) + O_n(i), \sigma_i(\theta)$), where $O_n(i)$ is the value obtained at the closest index to $i, n(i)$, which was already observed. This was found to empirically ease the learning task since, loosely, the network does not need to start “from scratch.”

Inputs to Q_i, K_i, V_i in the first attention layer: It is useful (empirically) to add the following features:

- empirical difference: $\Delta_i = O_i - O_{n(i)}$
- distance: $\Delta_{X_i} = t_i - t_{n(i)}$
- empirical derivative: $d_i = \frac{\Delta_i}{\Delta_{X_i}}$
- closest value: $O_{n(i)}$
- closest time-point: $t_{n(i)}$

So the inputs to the non-linear layers leading to K_i, V_i will be the concatenation of $(P'_i, \Delta_i, \Delta_{X_i}, d_i, O_{n(i)}, P'_{n(i)})$. and the inputs leading to Q_i will be the concatenation of $(P'_i, 0, \Delta_{X_i}, 0, O_{n(i)}, P'_{n(i)})$.

These two parts of the Taylorformer give the model a flavor of a Taylor approximation since we predict the current value as the closest value plus some function of the distance and derivative of closest values.

Additional attention mechanism. It is constructed in the Taylorformer (hereafter multihead attention X, or “MHA-X”). multiple nonlinear layers will be applied to P'_i leading to the two vectors Q_i^x and K_i^x . The vector V_i^x will be the result of applying a few nonlinear layers only on O_i . The “additional weighted attention” from this branch is calculated as $B_\theta^x \left(\sum_{j=1}^{i-1} A_{i,j}^x \times V_j^x \right)$. Here, $A_{i,j}^x$ is the dot-product of Q_i^x with K_j^x . The “additional weighted attention” from the first layer is then used as the updated values for the vectors Q_i^x, K_i^x . For V_i^x at each layer, the input stays the same as in the first layer. This attention mechanism has been shown to empirically help learning in the Taylorformer paper.

Connecting the two attention mechanisms. In the last layer, before prediction, we concatenate the outputs from the “classic” and additional attention mechanisms.

D.5. Inputs for a BC task using Taylorformer decoder

Here, the main adaptation needed is to take the two time-series, the observations and climate model outputs, and fit them into the 1-dimensional formulation presented above.

There is a notational convenience we can use here, which will uncover that what we refer to as the correction task is just a special case of prediction—once this becomes apparent, it is clear that all the standard ML tools for prediction can be translated into tools for correction. We define two-dimensional coordinates that encode both the time point and which time series we are looking at

$$[x_1, \dots, x_N] = [(t'_1, 2), \dots, (t'_n, 2), (t_1, 1), \dots, (t_m, 1)] \quad \text{where } N = m + n$$

and let the generalized time-series be $[y_1, \dots, y_N] = [g_1, \dots, g_n, o_1, \dots, o_m]$.

The task is then to predict $\hat{y}_1, \dots, \hat{y}_\ell$ at new coordinates $\hat{x}_1, \dots, \hat{x}_\ell = (\hat{t}_1, 1), \dots, (\hat{t}_\ell, 1)$. With the notation for a “generalized” time series, it is easy to insert it to the Transformer/Taylorformer machinery. Note that the additional features for the Taylorformer are calculated separately for the observations and for the climate model outputs since, for instance, we do not want to mix derivatives’ values between the two time series. For full details of the training and inference procedures for BC refer back to the main text.

E. Discussion: other ways to slice and infer the data

In Algorithm 1 (and the additional specification for multiple initial conditions), we have detailed a specific slicing procedure for the Taylorformer consisting of the steps “Initial condition selection”→ “Time window selection”→ “Prediction index selection”. But this is just one possibility. We now explain what is the reasoning behind a valid slicing operation.

E.1. Slicing discussion

It is not trivial to slice the pair of the full trajectories of (climate model, observations) sequences. In Figure 2, the first column shows the full sequences, and the second column shows example slices. The purpose of the slicing operation is to encourage the correction task instead of a prediction task: if we do not slice the data and set each $O^{*i} = [O_{t_1}, \dots, O_{t_n}]$ and $\mathbf{g}^{i,z} = [g_1^z, \dots, g_n^z]$ containing all climate model outputs just for different initial-condition indexed by z , then the ML model will learn to ignore the climate model outputs

completely; the observation value would always be the same for all pairs while the climate model output would keep on changing based on which initial-condition z was chosen and hence uninformative.

Given the understanding that we have to make our pairing (by some slicing) process in an educated manner, how should we pair the climate model(s) with observations? This is only possible if we make some assumptions about the generating process.

Here are a few possibilities for sensible corpus constructions:

1. Construct a dataset comprised of many locations and multiple initial conditions. Each row describes a unique location and is paired with the initial condition index z to give the sequence of (climate model z , observations) for all time points.
2. Construct a dataset of one unique location, multiple initial conditions, and multiple time-windows. Each row describes the same location and is paired with initial condition index z over changing unique time-windows to give the sequence of (climate model z , observations).

What makes a corpus construction sensible for our task is its ability to do corrections, i.e., use the information provided by the climate models as opposed to ignoring them and just using information based on observations and time points. Intuitively, for the first option, the location is changing from sequence to sequence, and thus a fixed “rule” that fits to location A using solely the observations will fail when applied on a different location B (assuming A and B are not very similar). For the second option, we have the same location for all sequences. But, a “fixed” rule will again fail since it has to be appropriate for different time windows. For example, if one sequence is from 1948 and another is from 1980, the same “rule” would not work unless the process is stationary (loosely, not changing its properties over time).

These configurations are selected examples of slicing the data. The question remains: which possibility should be chosen to construct the corpus? A standard ML procedure is to choose the model that performed best on hold-out likelihood and this is how we chose our slicing above.

E.2. Inference discussion

The Taylorformer architecture permits varying-length sequences. This flexibility is useful when training the model since we do not know what the optimal sequence length is when slicing the original sequence into subsequences. However, during inference, this flexibility poses a challenge—what should we feed into the trained model to generate the 20 years of values we are interested in? We could have used different setups to execute the sampling (generation) process, and not all generation processes are equal. It is a matter of empirical work to choose the generation process. Our choice above is a convenient choice to align loosely with the training procedure. Alternatively, we could inject the entire time series of a climate model run and all observations from 1948 to 1988 and sample day-by-day values up until 2008. Or, we can use the same inputs to generate in nonchronological order, generate the 1st of January 1989, and then fill in the values in-between before continuing to 1990 and further.

Table 1. Average MSE and Log-Likelihood for Abuja and Tokyo (1989–2008)

	MSE (Abuja)	Log-Likelihood (Abuja)	MSE (Tokyo)	Log-Likelihood (Tokyo)
IPSL (Lurton et al., 2020)	11.12	−2.62	39.57	−3.26
Meanshift (Xu, 1999)	5.86	−2.30	14.45	−2.75
mean + variance (Ho et al., 2012)	7.23	−2.41	12.29	−2.67
EQM (Panofsky & Brier, 1968)	7.69	−2.44	12.23	−2.67
dTsmbc (Robin & Vrac, 2021)	21.55	−2.95	13.28	−2.71
EC-BC (Vrac & Friederichs, 2015)	24.01	−3.01	129.41	−3.85
Temporal BC	3.57	−1.78	6.63	−2.41

Model Performance Comparison (1989–2008): This table presents the average Mean Squared Error (MSE) and log-likelihood for various models, based on 32 initial condition runs, assessing performance in Abuja and Tokyo. Lower MSE values indicate better performance, while higher log-likelihood values are preferable. Our temporal BC ML model is highlighted in gray. For models other than our Temporal BC, the probability distribution is assumed to be Normal, with mean $\mu_{i,model}$ (the model’s output at time t_i) and a constant variance σ_{model}^2 . Here, $\sigma_{model}^2 = \frac{1}{N} \sum_{i=1}^N (\mu_{i,model} - o_i)^2$.

F. Additional experimental results




Figure 19. QQ plots for Tokyo, Japan 1989–2008 for maximum daily temperatures (T_{max}) for observations (X -axis) versus BC methods (Y -axis): The climate model (IPSL) quantiles (top left plot) are very different than observations (ERA5). For mean shift (bottom left), EQM (bottom right), and our method **Temporal BC** (top right), the quantiles are captured quite well with minor differences.




Figure 20. QQ plots for Abuja, Nigeria 1989–2008 for maximum daily temperatures (T_{max}) for observations (X -axis) versus BC methods (Y -axis): The climate model (IPSL) quantiles (top left plot) are very different than observations. For mean shift (bottom left) and EQM (bottom right) the quantiles are captured more accurately but our method **Temporal BC** (top right) seems to be the most accurate.




Figure 21. Partial auto-correlation plot for Tokyo, Japan, 1989–2008 for lags 2–14 and for a randomly chosen initial condition run. Our Taylorformer temporal BC (horizontal blue lines) outputs multiple trajectories and, therefore, multiple partial auto-correlation values at each lag; it seems to underestimate the observed partial auto-correlation (horizontal orange line) for lags 2 and 3 and capture it well for higher lags. See the markers at the bottom of the figure to interpret the partial auto-correlation for other BC methods.




Figure 22. Partial auto-correlation plot for Abuja, Nigeria, 1989–2008 for lags 2–14 and for a randomly chosen initial condition run. Our Taylorformer temporal BC (horizontal blue lines) outputs multiple trajectories and, therefore, multiple partial auto-correlation values at each lag; it seems to capture the observed partial auto-correlation (horizontal orange line) for all lags except lag 3. See the markers at the bottom of the figure to interpret the partial auto-correlation for other BC methods.

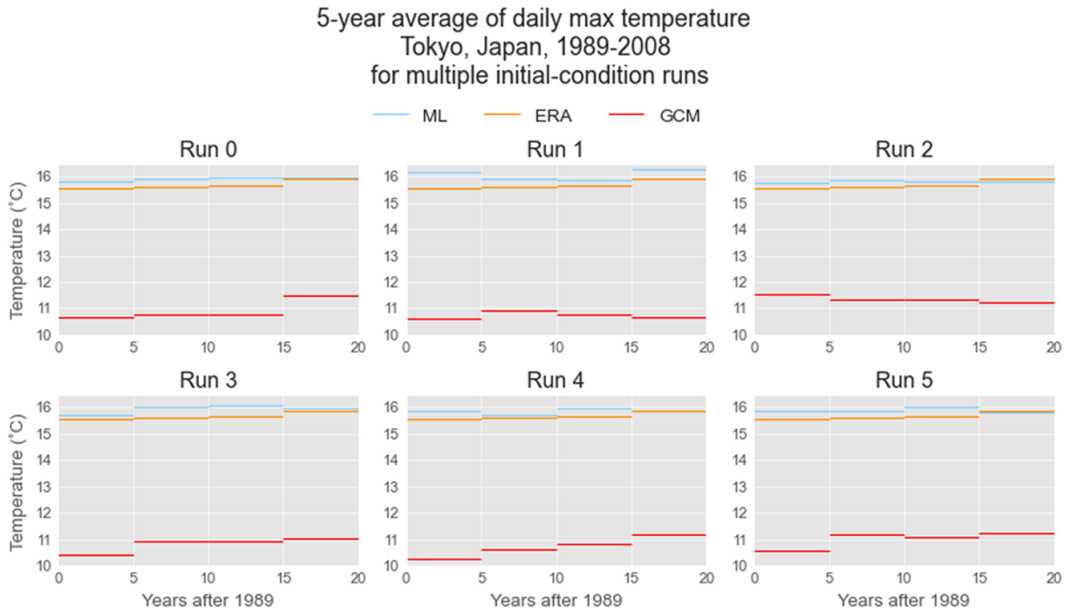


Figure 23. Comparison of trend evolution of Climate Model, Observations and Temporal BC in Tokyo, Japan (1989–2008): This figure illustrates the 5-year averages of daily maximum temperatures, with each plot representing a different initial-condition run. The Temporal BC model is depicted by a blue horizontal line, which does not necessarily align with the trend of the initial-condition climate model, shown as a red horizontal line. Furthermore, neither model consistently follows the trend observed in the actual data, represented by an orange horizontal line.

F.1. Tokyo, Japan

Figure 25 shows the results for Tokyo, Japan using a heatwave threshold of three consecutive days above 24°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989–2008. Our novel Taylorformer temporal BC produces a more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 0.7% from the observed number of heatwaves. Most other BC models perform worse: 12.3%,4.2%,4.1%,35.6%,89%,4.2% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and Tsmbc, respectively.

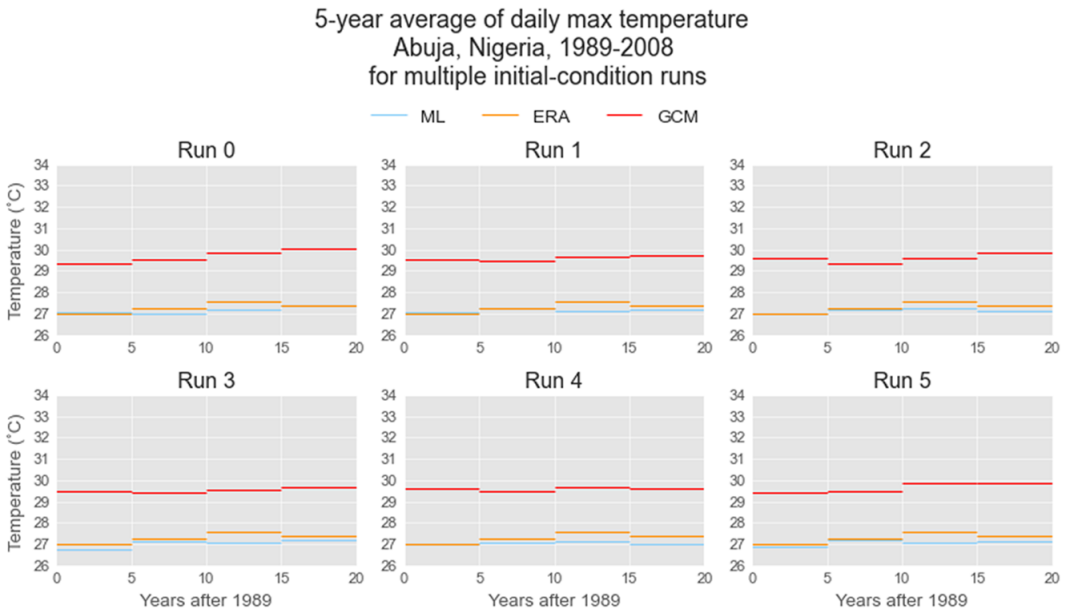


Figure 24. Comparison of trend evolution of Climate Model, Observations and Temporal BC in Abuja, Nigeria (1989–2008): This figure illustrates the 5-year averages of daily maximum temperatures, with each plot representing a different initial-condition run. The Temporal BC model is depicted by a blue horizontal line, which does not necessarily align with the trend of the initial-condition climate model, shown as a red horizontal line. Furthermore, neither model consistently follows the trend observed in the actual data, represented by an orange horizontal line.




Figure 25. Comparative Analysis of “number of heatwaves” Trends in Tokyo, Japan (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 24°C is shown. The IPSL climate model predictions are represented by red triangles, which generally underestimate the observations. Actual observations are indicated by a vertical orange line. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

Figure 26 shows the results for Tokyo, Japan using a heatwave threshold of three consecutive days above 26°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989–2008. Our novel Taylorformer temporal BC produces a more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 5.6% from the observed number of heatwaves. Most other BC models perform worse: 15%, 11%, 7%, 56%, 99%, 5.5% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.




Figure 26. Comparative analysis of “number of heatwaves” trends in Tokyo, Japan (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 26°C is shown. The IPSL climate model predictions are represented by red triangles, which generally underestimate the observations. A vertical orange line indicates actual observations. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

Figure 27 shows the results for Tokyo, Japan using a heatwave threshold of three consecutive days above 28°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989 – 2008. Our novel Taylorformer temporal BC produces a much more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 18.8% from the observed number of heatwaves. Other BC models perform worse: 75.3%,42.8%,40%,94.7%,99%,47.7% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and Tsmbc, respectively.




Figure 27. Comparative analysis of “number of heatwaves” Trends in Tokyo, Japan (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 22°C is shown. The IPSL climate model predictions are represented by red triangles, which generally underestimate the observations. A vertical orange line indicates actual observations. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

F.2. Abuja, Nigeria

Figure 28 shows the results for Abuja, Nigeria, using a heatwave threshold of three consecutive days above 26°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989–2008. Our novel Taylorformer temporal BC produces a more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 4.9% from the observed number of heatwaves. Most other BC models perform worse: 3.7%, 11%, 14.8%, 42.9%, 89.8%, 65.9% for mean shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.




Figure 28. Comparative Analysis of “number of heatwaves” Trends in Abuja, Nigeria (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 26°C is shown. The IPSL climate model predictions are represented by red triangles, which generally overestimate the observations. A vertical orange line indicates actual observations. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

Figure 29 shows the results for Abuja, Nigeria using a heatwave threshold of three consecutive days above 28°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989 – 2008. Our novel Taylorformer temporal BC produces a much more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 10.2% from the observed number of heatwaves. Most other BC models perform worse: 6.5%,9.3%,14.5%,59.3%,91.9%,71.9% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.




Figure 29. Comparative Analysis of “number of heatwaves” Trends in Abuja, Nigeria (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 28°C is shown. The IPSL climate model predictions are represented by red triangles, which generally overestimate the observations. A vertical orange line indicates actual observations. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

Figure 30 shows the results for Abuja, Nigeria using a heatwave threshold of three consecutive days above 30°C maximum temperature. The raw outputs from the IPSL climate model (red triangles) underestimate the number of instances of observed heatwaves (vertical orange line) for the period 1989 – 2008. Our novel Taylorformer temporal BC produces a much more accurate distribution (horizontal box-plots) per run (0–31) for the number of heatwaves in the same period differing on average by 23.6% from the observed number of heatwaves. Most other BC models perform worse: 21.7%,10.3%,13.5%,73.7%, 136%,67.6% for mean shift, mean and variance shift, EQM, EC-BC, 3D-BC, and TSMBC, respectively.




Figure 30. Comparative Analysis of “number of heatwaves” Trends in Abuja, Nigeria (1989–2008): The number of periods featuring at least three consecutive days with temperatures exceeding 30°C is shown. The IPSL climate model predictions are represented by red triangles, which generally overestimate the observations. A vertical orange line indicates actual observations. The Taylorformer temporal BC is depicted using horizontal box plots, with whiskers indicating the 1st and 3rd quartiles. Markers for other BC methods are indicated at the bottom of the figure.

G. Supplementary figures




Figure 31. When collapsing a time-series (top) to a histogram (bottom), the temporal information is lost – a toy illustration.

H. Gaussian Process models

Gaussian process models used. For a time series with random variables Y_1, \dots, Y_N for time-points t_1, \dots, t_N the GP in Figure 16 can be written as

$$Y_i = Z_i$$

where Z_i is the i -th element from the vector $Z \sim \text{MVN}(0, K)$, with $K_{ij} = \exp\left(-\frac{(t_i - t_j)^2}{2\ell^2}\right)$ (this is also known as an RBF kernel) ℓ is called the length-scale and for the figure we took $\ell = 0.1$.

The GP in Figure 17 can be written as

$$Y_i = Z_i * W_i$$

where Z_i is the i -th element from the vector $Z \sim \text{MVN}(0, K)$, with $K_{ij} = \exp\left(-0.5 * \frac{(\sin(\pi r) / \gamma)^2}{\ell^2}\right)$ (this is also known as a Periodic kernel) r, ℓ, γ is the euclidean distance between t_i and t_j , ℓ is the length-scale and γ is the period parameter, all set to one. W_i is the i -th element from the vector $W \sim \text{MVN}(0, K)$, with $K_{ij} = \left(\frac{1 + r^2}{2 * \alpha * \ell^2}\right)^{-\alpha}$ (this is also known as a Rational Quadratic kernel). Here, α determines the relative weighting of small-scale and large-scale fluctuations.

The GP in Figure 18 can be written as

$$Y_{i,k} = Z_i + E_k$$

where Z_i is the i -th element from the vector $Z \sim \text{MVN}(0, K)$, with $K_{ij} = \exp\left(-\frac{(t_i - t_j)^2}{2\ell^2}\right)$ (this is also known as an RBF kernel). $E_k \sim \text{Normal}(0, 2)$, with $k = 1$ representing “observations” and $k = 2$ representing “climate model” outputs. Note that the time mismatch is created by shifting “observations” timestamps by 0.9 so that for $k = 2$ we have time points $\mathbf{t} = (t_1, \dots, t_N)$ and for $k = 1$ we have $\mathbf{t} + 0.9$.

I. BC models—implementation details

We have compared our Taylorformer to six BC models. We detail their implementations below:

Mean shift. Calculate the difference in means between the climate model and observations for each month and initial condition run over the reference period (1948 – 1988). Then, apply the monthly difference to the climate model in the same month during the projection period (1989 – 2008).

Mean and variance shift. For each month and climate model initial condition run over the reference period (1948 – 1988), calculate the means and standard deviations (std) of the climate model and separately for observations. Then, for each climate model value in the same month over the projection period (1989 – 2008), subtract the climate model mean of the reference period, multiply by the ratio of observation std to climate model std and finally add the observed mean during the reference period. (Our implementation follows Hawkins et al. (Hawkins et al., 2013).)

Empirical Quantile Mapping. For each month and climate model initial condition run over the reference period (1948 – 1988), sort the observations and separately the climate model outputs. For the unsorted values during the same month during the projection period, find the index of the first value bigger or equal in the sorted climate model data (during reference). The corrected value is then the observed value at that index.

EC-BC. First, perform EQM as detailed above. Then, for each climate model initial condition run, get the ranks of observations during the reference period (1968 – 1988). Use these ranks to resort the climate model outputs during the projection period (1989 – 2008).

3D-BC. For each climate model initial condition run, we follow the exact algorithm provided in the paper by Mehrotra et al. (Mehrotra & Sharma, 2019). The reference period is 1968 – 1988, and the projection period is 1989–2008.

TSMBC. Here, we have used the R package “SBCK” (Robin & Vrac, 2023). We initialized it with the code: `dTSMBC$new(lag = 10)`. Then, for each climate model initial condition run, we used the code `tsmbc$fit(Y0, X0, X1)` followed by `tsmbc$predict(X1)`, where *Y0* stands for observations during the reference period, and *X0, X1* for climate model outputs in the reference and projection period, respectively.