Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Machine Learning Models for Subnational Climate Forecasting and Forest Ecosystem Monitoring across the 34 Prefectures of Guinea in West Africa

Kossi Kuma Katakpe

kossi.katakpe@imsp-uac.org

University of Labé

Adoté Hervé Gildas AKUESON

Institut Sup´erieur Agronomique et V´et´erinaire Valery Giscard d'Estaing de Faranah

Moustapha N'gonet Camara

University of Labé

Abdoulaye Sow

University of Labé

Pépé Kpoghomou

University of Labé

Mohamed Makanera

University of Labé

Research Article

Keywords: Climate forecasting, Multilayer perceptron, Random forest, Gradient boosting, West African Monsoon, Forest-ecosystem monitoring, Guinea, Extrapolation, Machine learning

Posted Date: June 22nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10084122/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Machine Learning Models for Subnational Climate Forecasting and Forest Ecosystem Monitoring across the 34 Prefectures of Guinea in West Africa

Kossi Kuma Katakpe[1,2*], Adoté Hervé Gildas AKUESON[3,4], Moustapha N'gonet Camara[1], Abdoulaye Sow[1], Pépé Kpoghomou[1], Mohamed Makanera[1]

[1*]Department of Computer Science, UL – University of Labé, Labé, Guinea.

[2]Department of Mathematics and Computer Science, Institute of Mathematics and Physical Sciences (IMSP), Porto-Novo, Benin.

[3]Department of Water, Forests and Environment, Laboratoire de Foresterie, Environnement et Modélisation Écologique (LaFEME), Unité de Biométrie, Modélisation et Inventaires Forestiers, Institut Supérieur Agronomique et Vétérinaire Valery Giscard d'Estaing de Faranah, BP 131 Faranah, Republic of Guinea.

[4]Applied Statistics and Computer Science Unit (USIA), Laboratory of Forestry Studies and Research (LERF), Faculty of Agronomy, University of Parakou, P.O. Box 123, Parakou, Republic of Benin.

*Corresponding author(s). E-mail(s): kossi.katakpe@imsp-uac.org;

Abstract

Reliable sub-national climate forecasts are a prerequisite for monitoring the forest ecosystems of West Africa, where rainfall and temperature regimes govern phenology, fire risk, species distribution and the livelihoods that depend on forest resources. A recent study modelled monthly temperature and precipitation across all 34 prefectures of Guinea with single-layer (simple) perceptrons and explicitly identified the linear architecture as the principal barrier to higher accuracy. Here we revisit the same prefecture-level problem with a substantially larger record (1 020 monthly values per prefecture, 2015–2099) and conduct a controlled comparison of four model classes under identical inputs and an identical chronological

train/validation/test protocol: the simple perceptron baseline, a multilayer perceptron (MLP) neural network, and two ensemble-learning methods (random forest and gradient boosting). Across 544 model fits we find a clear, physically interpretable dissociation between the two families of nonlinear learners. For temperature, which is dominated by a strong seasonal cycle and a pronounced warming trend, the MLP is the best and most spatially robust model (national mean $\boldsymbol{R^2 = 0.86}$), whereas the tree ensembles fail systematically ($\boldsymbol{R^2 = 0.59}$ for random forest, with negative $\boldsymbol{R^2}$ in four coastal prefectures) because they cannot extrapolate beyond the range of values seen during training. For precipitation, a bounded, highly skewed and quasi-stationary variable, the ranking reverses: random forest is the strongest model ($\boldsymbol{R^2 = 0.89}$, a 24% reduction in RMSE relative to the perceptron), while the MLP offers no advantage. The neural-network advantage for temperature widens at longer lead times (12-month $\boldsymbol{R^2 = 0.76}$ for the MLP versus 0.69 for the perceptron and 0.45 for random forest) and is largest precisely in the Forest-Guinea and coastal prefectures most relevant to ecosystem monitoring. We translate these results into a practical, variable-specific forecasting strategy for an operational forest-monitoring system and argue that the inability of ensemble methods to extrapolate is a decisive, and frequently overlooked, consideration for any climate-change application in which the future lies outside the historical envelope.

Keywords: Climate forecasting, Multilayer perceptron, Random forest, Gradient boosting, West African Monsoon, Forest-ecosystem monitoring, Guinea, Extrapolation, Machine learning

1 Introduction

Climate change is reshaping West Africa, and few regions are as exposed: its forests, its farming and its water supply all respond sharply to even modest shifts in temperature and rainfall [1]. Guinea is a useful place to see this play out. Often called the "water tower of West Africa", the country packs several climates into a small area: a narrow Atlantic littoral, the orographic massif of the Fouta Djallon, the continental interior of Upper Guinea, and the sub-equatorial forest block of the south, each with its own thermal and pluviometric signature. Where dense humid forest, gallery forest, woodland savanna or montane vegetation can grow depends directly on these regimes [2–5], which also set the phenological calendars, the fire seasons, and the hydrological services that forest ecosystems and the people who depend on them rely upon. For anyone trying to monitor those forests, plan restoration or run an early-warning system, then, spatially resolved climate forecasts are less a research nicety than a working necessity [6].

Process-based general circulation and Earth-system models remain the foundation of long-range climate projection [1, 7], but their coarse spatial resolution, heavy data and computing requirements, and reliance on sub-grid parameterisations limit their direct usefulness at the prefecture scale, particularly in data-sparse African settings [8]. Long-term reanalysis products such as ERA5 [9] have partially alleviated this observational scarcity, yet they cannot replace the spatial detail that statistical

downscaling must provide [10]. This has motivated a rapid expansion of data-driven approaches. Artificial neural networks, in particular, are universal function approximators [11, 12] able to learn nonlinear relationships directly from observations without an explicit physical formulation, and a large literature now documents their value for temperature and rainfall prediction, statistical downscaling and bias correction [13–16]. Recent transformer- and graph-based architectures have achieved unprecedented skill at global medium-range weather forecasting [17, 18], confirming the central role of neural networks in modern Earth system modelling [19]. In parallel, ensemble-learning methods based on regression trees (random forests [20] and gradient boosting [21]) have become the default choice for tabular prediction in many applied sciences because of their accuracy, robustness to noise and minimal need for feature scaling.

A recent contribution modelled monthly temperature and precipitation across all 34 prefectures of Guinea using single-layer (simple) perceptrons and reported a useful but clearly bounded level of skill, with national mean coefficients of determination of approximately 0.55 for temperature and 0.67 for precipitation over a 2015–2024 record [22]. That study was explicit about the cause of its ceiling: the linear, single-layer architecture of the perceptron cannot represent the nonlinear interactions, threshold behaviour and skewed distributions that characterise tropical climate, and it identified multilayer networks and more expressive learners as the natural next step. The present work takes up that challenge directly.

We could simply ask whether a deeper network beats the perceptron, but that is not the question a practitioner actually faces. The real question is a choice: when a forest-monitoring agency has to pick a learner for sub-national climate forecasting, is it better off with a neural network or with an ensemble method, and does the answer change depending on whether it is temperature or rainfall that is being predicted? To settle this we use a much larger dataset, 1 020 monthly values per prefecture over 2015–2099, and we put four model classes (the simple perceptron, a multilayer perceptron, a random forest and a gradient-boosting machine) through exactly the same inputs, preprocessing and chronological evaluation. The answer turns out to be a clean split. Neural networks win for temperature; ensembles win for precipitation. What drives the split is the statistical character of each variable and, above all, whether a model can extrapolate the warming trend, something neural networks do naturally and tree-based models cannot. The rest of the paper works through what this means for operational forest-ecosystem monitoring in Guinea and across West Africa more broadly.

2 Related work

2.1 Classical climate modelling and its limits

Classical approaches to climate prediction fall into statistical, dynamical and empirical families. Statistical models exploit historical relationships among variables and provide computationally cheap short-range forecasts, but their accuracy degrades when variability is high or when the relationships are nonlinear [13]. Dynamical models solve the governing equations of the ocean–atmosphere system and are indispensable for understanding large-scale phenomena [1, 7], yet they are expensive to calibrate and run, and their coarse resolution leaves key local processes (convection, orographic

rainfall, land–atmosphere feedbacks) unresolved [7]. Empirical models are simple and transparent but struggle with extremes and long-term trends [6]. Across all three families, the explicit representation of small-scale, nonlinear processes remains a central difficulty, which is precisely the gap that machine learning is increasingly used to fill, either as a stand-alone predictor or as a hybrid component within physical models [23–26]. Initialized prediction systems now extend this to subseasonal-to-decadal timescales [27].

2.2 Neural networks for climate prediction

Artificial neural networks are layered compositions of weighted connections and nonlinear activation functions trained, typically by back-propagation [28], to minimise a prediction error. The perceptron introduced by Rosenblatt [29] provided the conceptual seed, but its single-layer form is restricted to linearly separable problems, as Minsky and Papert famously showed [30]. The addition of one or more hidden layers yields the multilayer perceptron, which the universal approximation theorems of Cybenko [11], Funahashi [31] and of Hornik, Stinchcombe and White [12] established as capable of approximating any continuous function to arbitrary accuracy. In climate science, neural networks have been applied successfully to rainfall estimation, surface-temperature forecasting and the modelling of large-scale climate indices [14, 15, 32], and they are well suited to the high-dimensional, non-stationary datasets typical of meteorology [16, 33–35]. At the global frontier, data-driven architectures such as Pangu-Weather [17] and GraphCast [18] now match or surpass operational numerical weather prediction, while foundation models for weather and climate [36] open further perspectives. Physical constraints can be enforced in neural network designs to improve thermodynamic consistency [37, 38]. Within West Africa, recurrent and feed-forward architectures have delivered strong seasonal-rainfall forecasts in Ghana [39, 40], Nigeria [41] and Benin [42], repeatedly outperforming classical regression in data-scarce settings [43], while hybrid convolutional–recurrent models have advanced multi-variable monthly prediction at larger scales [44, 45].

2.3 Ensemble learning

Ensemble methods combine many weak learners into a single strong predictor. The random forest of Breiman [20] aggregates de-correlated regression trees grown on bootstrap samples with randomised feature selection, achieving low variance and strong out-of-the-box accuracy. Gradient boosting, formalised by Friedman [21], instead builds trees sequentially, each correcting the residual errors of its predecessors, and its modern implementations including XGBoost [46] dominate many predictive-modelling benchmarks. Both methods are attractive for environmental tabular data because they capture nonlinear interactions and threshold effects automatically and require no feature scaling [47, 48]. Side-by-side comparisons of several such algorithms on a common task have become routine across applied fields, from environmental modelling to computer vision [43, 49], which is precisely the spirit of the benchmark we pursue here. A structural property of all tree-based learners, however, is that their predictions are piece-wise constant averages of training targets and therefore cannot exceed the range

of values observed during training. This limitation is benign for stationary data but becomes critical whenever the prediction target drifts beyond the historical envelope, as it does, by construction, under a warming trend [24, 50]. Comprehensive climate-projection benchmarks confirm that this structural constraint systematically degrades tree-ensemble performance on trended targets [51].

2.4 Positioning of the present study

Previous work on Guinea established a transparent linear baseline and called for more expressive models [22]. We answer that call, but we go further by framing the problem as a head-to-head comparison between the two leading families of nonlinear learners, neural networks and tree ensembles, under identical conditions and over a record long enough to expose the consequences of a warming trend. To our knowledge this is the first prefecture-resolved, multi-model benchmark of this kind for Guinea, and the first to connect the extrapolation behaviour of the competing methods explicitly to the needs of forest-ecosystem monitoring. Neural networks' ability to detect physically meaningful patterns of forced climate change directly from data has been demonstrated in several recent attribution studies [52, 53], lending further support to their use in a warming-trend context.

3 Study area and data

3.1 Study area

Guinea (approximately 245 860 km^2) is administratively divided into 34 prefectures grouped into four natural regions that coincide closely with its principal climatic and vegetation zones. Maritime (Lower) Guinea, comprising the nine coastal prefectures, experiences an oceanic, high-rainfall regime with damped thermal amplitude. The Fouta Djallon highlands of Middle Guinea (ten prefectures) generate intense orographic precipitation and a moderate, elevation-regulated temperature regime; their montane forests and headwaters feed the major rivers of West Africa. Upper Guinea (eight prefectures) has a continental Sudano-Guinean climate with the country's widest seasonal temperature range and a drier, woodland-savanna landscape. Forest (south-eastern) Guinea (seven prefectures) has a sub-equatorial climate with abundant, more evenly distributed rainfall that sustains the country's densest humid forests. This regional diversity makes Guinea an ideal natural laboratory for evaluating how model performance varies with climatic regime, and it maps directly onto the forest types whose monitoring motivates the study.

3.2 Data

Two prefecture-level datasets were used: monthly mean temperature (°C) and monthly total precipitation (mm) for each of the 34 prefectures over the period January 2015 to December 2099, giving 1 020 monthly observations per prefecture and variable. Annual aggregates of the same series were used to characterise long-term tendencies. Over this record the prefecture-averaged temperature rises by approximately 0.71 °C per decade (statistically significant), amounting to a mean warming of about 5.34 °C

between the 2015–2024 reference period and 2090 - 2099; every prefecture warms, with decadal rates between roughly 0.53 and 0.78 °C. Precipitation shows no significant national trend (about −0.63 mm per decade) but pronounced and spatially contrasted prefecture-level changes. The presence of a strong, near-monotonic warming signal is central to the analysis that follows, because it places the latter part of the record outside the range of the earlier part, the regime in which extrapolation ability becomes decisive. These warming rates are consistent with recent multi-model assessments of future temperature change over West Africa [54, 55], and the precipitation series are consistent with CMIP6-based projections validated over the West African monsoon region [56, 57]. The ERA5 reanalysis [9] provided the historical observational anchor for the 2015–2024 calibration period.

Compared with the 2015–2024 window of the earlier perceptron study [22], the full 2015–2099 record increases the number of monthly observations per prefecture roughly twelve-fold. This larger sample directly addresses the small-data limitation noted previously and provides the volume of training data that more flexible learners require to realise their potential.

4 Methods

4.1 Forecasting problem and features

We treat climate forecasting as a supervised one-step-ahead regression problem solved independently for each prefecture and each target variable. To predict the value of a variable in month t we use the twelve preceding monthly values of that variable as autoregressive lag features; a full year of lags allows the models to recover the seasonal cycle. To this we add a cyclical encoding of the calendar month and a normalised time index $\tau_t = t/T \in [0, 1]$ that supplies an explicit, monotonic representation of long-term tendency. Writing y_t for the target in month t and $m_t \in \{1, \ldots, 12\}$ for its calendar month, the univariate predictor vector is

$$\mathbf{x}_t = \Big[\, y_{t-1},\, y_{t-2},\, \ldots,\, y_{t-12},\, \sin\!\big(\tfrac{2\pi m_t}{12}\big),\, \cos\!\big(\tfrac{2\pi m_t}{12}\big),\, \tau_t \,\Big]^{\top}. \tag{1}$$

The sine cosine pair encodes the month on a circle so that December and January sit next to each other, while τ_t lets a model read off where in the long warming record a given month falls. In the univariate scenario the predictors comprise the target's own lags plus these three calendar/time features (15 features); in the multivariate scenario the twelve lags of the second variable are added (27 features), allowing the models to exploit temperature precipitation coupling. Crucially, all four models receive exactly the same inputs, so that any difference in performance reflects the model class rather than the feature engineering.

Because monthly precipitation is strongly right-skewed, with many near-zero dry-season values and large wet-season peaks, the precipitation target and its lagged predictors were log-transformed using $\log(1 + x)$ before modelling and back-transformed to millimetres before all error metrics were computed. The same transformation was applied to every model, preserving the fairness of the comparison.

4.2 Models

All four learners map the same predictor vector $\mathbf{x}$ of Eq. (1) to a one-step-ahead estimate $\hat{y}$, but they build that map in very different ways. Linear and neural models received z-score standardised inputs $\mathbf{z} = (\mathbf{x} - \boldsymbol{\mu})/\boldsymbol{\sigma}$, with the mean $\boldsymbol{\mu}$ and standard deviation $\boldsymbol{\sigma}$ estimated on the training set only; the tree models are scale-invariant and used the raw features. Figure 1 summarises the four architectures and the shared pipeline. We give the governing equations below because they make the central contrast of this paper (which models can follow a trend out of the training range and which cannot) explicit rather than merely verbal.

The **simple perceptron** is the linear, single-layer baseline of the earlier study [22]. It forms a weighted sum of the standardised inputs,

$$\hat{y}(\mathbf{x}) = \mathbf{w}^{\top}\mathbf{z} + b, \tag{2}$$

and its weights are obtained by minimising a ridge-penalised squared error,

$$(\mathbf{w}^{\star}, b^{\star}) = \arg\min_{\mathbf{w},\, b} \sum_{i=1}^{N} \big(y_i - \mathbf{w}^{\top}\mathbf{z}_i - b\big)^2 + \lambda \,\|\mathbf{w}\|_2^2, \tag{3}$$

which is the closed-form optimum reached by stochastic gradient descent on the mean-squared error with light L2 regularisation.

The **multilayer perceptron (MLP)** stacks two hidden layers of 64 and 32 rectified-linear units between the input and a single linear output. With $\phi(u) = \max(0, u)$ applied element-wise, a forward pass is

$$\mathbf{a}^{(1)} = \phi\big(W^{(1)}\mathbf{z} + \mathbf{b}^{(1)}\big), \quad \mathbf{a}^{(2)} = \phi\big(W^{(2)}\mathbf{a}^{(1)} + \mathbf{b}^{(2)}\big), \quad \hat{y} = \mathbf{w}^{(3)\top}\mathbf{a}^{(2)} + b^{(3)}. \tag{4}$$

The parameters $\theta = \{W^{(1)}, W^{(2)}, \mathbf{w}^{(3)}, \mathbf{b}^{(1)}, \mathbf{b}^{(2)}, b^{(3)}\}$ minimise a regularised mean-squared error,

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \big(y_i - \hat{y}_i\big)^2 + \frac{\eta}{2}\Big(\|W^{(1)}\|_F^2 + \|W^{(2)}\|_F^2\Big), \tag{5}$$

trained by back-propagation [28] with the Adam optimiser [58] (learning rate 10^{-3}, weight decay $\eta = 10^{-4}$, mini-batches of 64) and early stopping on the validation set. To damp the sensitivity to weight initialisation we averaged the predictions of three independently initialised networks. Because the output layer is an unbounded affine combination of its inputs, $\hat{y}$ in Eq. (4) is free to take values never seen in training, the property that lets the MLP follow a warming trend.

The **random forest** [20] works on a completely different principle. A single regression tree partitions the feature space into L disjoint leaves $R_1, \dots, R_L$ and predicts,

for each leaf, the mean of the training targets that fall inside it,

$$T(\mathbf{x}) = \sum_{l=1}^{L} c_l \mathbb{1}\{\mathbf{x} \in R_l\}, \qquad c_l = \frac{1}{|R_l|} \sum_{i:\, \mathbf{x}_i \in R_l} y_i. \tag{6}$$

The forest averages B such trees, each grown on a bootstrap resample with a random subset of features tried at every split,

$$\hat{y}_{\mathrm{RF}}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^{B} T_b(\mathbf{x}). \tag{7}$$

Our forest uses $B = 50$ trees of maximum depth 10 with 45% of the features considered at each split. Because every leaf value c_l is an *average* of observed targets, each tree (and hence their mean) is bounded by the training range,

$$\min_i y_i \;\le\; \hat{y}_{\mathrm{RF}}(\mathbf{x}) \;\le\; \max_i y_i \qquad \text{for every } \mathbf{x}. \tag{8}$$

Inequality (8) is not a tuning choice but a structural ceiling, and it is the single fact that explains the temperature results of Section 5.2.

The **gradient-boosting machine** [21, 46] also combines regression trees, but sequentially rather than in parallel. Starting from the mean target $F_0(\mathbf{x}) = \bar{y}$, each stage adds a shallow tree h_m fitted to the pseudo-residuals of the squared-error loss,

$$r_{im} = -\left[\frac{\partial L\big(y_i, F(\mathbf{x}_i)\big)}{\partial F(\mathbf{x}_i)}\right]_{F = F_{m-1}} = y_i - F_{m-1}(\mathbf{x}_i), \tag{9}$$

and the model is updated with a learning-rate (shrinkage) factor ν,

$$F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \nu\, h_m(\mathbf{x}). \tag{10}$$

We use $M = 100$ depth-3 trees with $\nu = 0.08$ and 0.8 row sub-sampling. Since every h_m is itself a tree built from leaf averages, the boosted predictor inherits the same bounded behaviour as Eq. (8) and is, like the random forest, unable to extrapolate beyond the observed range. All four learners were implemented from first principles in Python/NumPy, which guarantees full transparency and exact reproducibility of every reported number.

4.3 Evaluation protocol

For each prefecture and variable the 1 008 lagged samples were split chronologically into a training set (the first 70%), a validation set (the next 10%, used for early stopping and model selection) and a held-out test set (the final 20%, corresponding to March 2083 to December 2099). The strictly chronological ordering prevents any leakage of future information into training and means that the test period lies almost

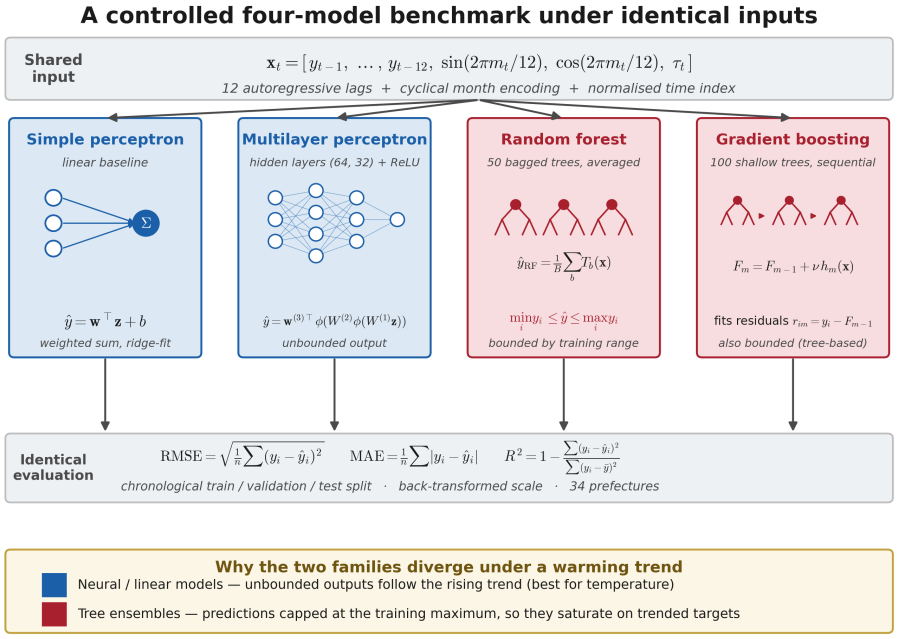


Fig. 1 Schematic of the four learners and the shared forecasting pipeline. All models receive the same predictor vector (Eq. 1) and are scored with the same metrics (Eq. 11) on an identical chronological split. The neural and linear family (blue) produces unbounded outputs and can follow the warming trend; the tree family (red) predicts averages of training targets and is therefore capped by the training range (Eq. 8), which is why the two families excel on different variables.

entirely beyond the temperature range of the training period a realistic and demanding setting for a climate-change application. Performance was quantified on the original (back-transformed) scale with three complementary metrics, the root-mean-square error (RMSE), the mean absolute error (MAE) and the coefficient of determination (R^2):

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}, \quad \mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}\left|\, y_i - \hat{y}_i \,\right|, \quad R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)^2}. \tag{11}$$

RMSE and MAE are reported in physical units (°C or mm) and R^2 is dimensionless; a model that merely predicts the test-period mean scores $R^2 = 0$, so a negative value signals a prediction worse than that trivial baseline. The full design comprises four scenarios (temperature and precipitation, each univariate and multivariate) across 34 prefectures and four models, i.e. 544 model fits. A complementary experiment evaluated direct forecasting at a 12-month horizon for temperature to probe the stability of each method at the seasonal-to-annual lead times most relevant to forest planning.

5 Results

5.1 Overall model comparison

Table 1 and Fig. 2 summarise national-mean performance in the univariate scenario. The results reveal a striking, variable-dependent dissociation between the two families of nonlinear learners. For temperature, the multilayer perceptron is the most accurate model ($R^2 = 0.86$, RMSE $= 0.79\,^{\circ}$C), marginally ahead of the simple perceptron ($R^2 = 0.85$) and far ahead of both ensemble methods, whose national skill collapses to $R^2 = 0.59$ for the random forest and $R^2 = 0.67$ for gradient boosting. For precipitation the ranking is reversed: the random forest is clearly the best model ($R^2 = 0.89$, RMSE $= 60.72$ mm), a 24% reduction in RMSE relative to the simple perceptron ($R^2 = 0.82$, RMSE $= 79.78$ mm), while the MLP trails both the ensemble methods and the linear baseline ($R^2 = 0.79$).

Two other numbers in Table 1 are worth dwelling on. The spread of skill across prefectures, captured by the standard deviation of R^2, is much wider for the tree ensembles on temperature (0.32 for the random forest) than for the neural and linear models (0.08 and 0.11). So the ensembles are not just less accurate on temperature on average; they are also a good deal less dependable from one prefecture to the next. Precipitation tells a calmer story: here all four models stay fairly even across space (R^2 standard deviations of just 0.03–0.05), which fits the broadly coherent monsoon signal.

Table 1 National-mean test-set performance across all 34 prefectures (univariate, one-step-ahead). Best R^2 per variable in bold. Temperature errors in °C, precipitation errors in mm.

Variable	Model	RMSE	MAE	R^2	R^2 SD
Temperature (°C)	Simple Perceptron	0.79	0.62	0.849	0.111
	MLP	0.79	0.62	**0.860**	0.082
	Random Forest	1.29	0.98	0.590	0.316
	Gradient Boosting	1.16	0.89	0.669	0.246
Precipitation (mm)	Simple Perceptron	79.78	44.24	0.820	0.032
	MLP	85.32	47.28	0.794	0.050
	Random Forest	60.72	34.39	**0.893**	0.031
	Gradient Boosting	63.41	35.94	0.884	0.032

5.2 Why ensembles fail on temperature: the extrapolation problem

The collapse of the tree ensembles on temperature is not a tuning artefact but a structural consequence of how tree-based models generate predictions [20]. Because the test period (2083–2099) is warmer than any part of the training period, the temperature values to be predicted, and the lagged features that encode them, lie largely outside the training range. By the bound of Eq. (8), a tree ensemble can never return a value

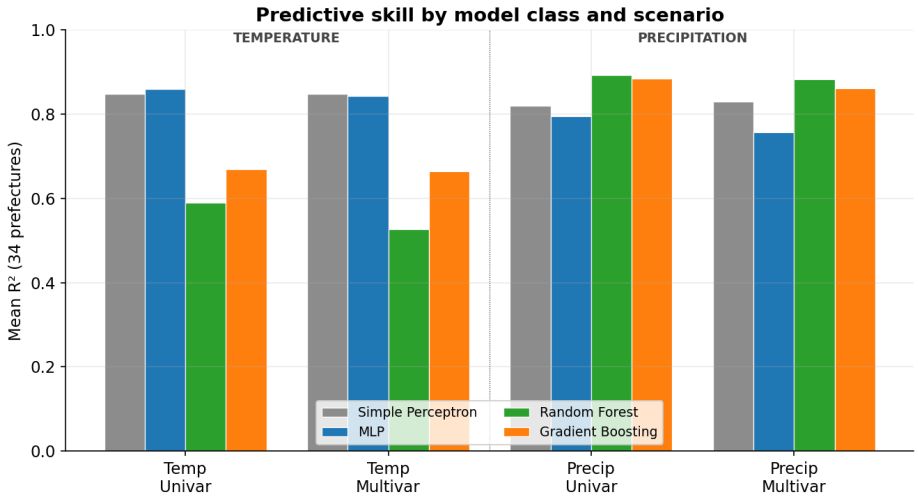


Fig. 2 Mean coefficient of determination (R^2) by model class and scenario across the 34 prefectures. Neural and linear models dominate temperature; ensemble methods dominate precipitation. The ensembles' temperature deficit reflects their inability to extrapolate the warming trend (Section 5.2).

above the warmest target it saw in training: its predictions saturate at that ceiling and simply cannot follow the continuing rise. Fig. 3 illustrates this directly for Siguiri, a continental Upper-Guinea prefecture with a strong warming trend: the random-forest predictions form a flat ceiling near 34.5 °C for all higher observed temperatures, whereas the MLP tracks the rising trend up to 38 °C. The effect is most severe in the coastal prefectures, where the seasonal temperature amplitude is small relative to the trend: the random forest returns negative R^2 in Boffa, Fria, Conakry and Dubreka, meaning it performs worse than simply predicting the test-period mean. Neural networks and linear models, whose outputs are unbounded linear combinations of their inputs, are free of this constraint and extrapolate the trend naturally [12, 28].

5.3 Per-prefecture skill gains over the baseline

Fig. 4 compares, prefecture by prefecture, the best neural model for temperature (MLP) and the best ensemble model for precipitation (random forest) against the simple-perceptron baseline. For temperature, the MLP lies on or above the 1:1 line in the great majority of prefectures, with the largest gains concentrated in the coastal cluster of Maritime Guinea, where high-frequency, sea-surface-temperature-driven variability is nonlinear and poorly captured by a linear model. For precipitation, the random forest lies above the diagonal in every single prefecture, confirming a uniform and substantial improvement over the linear baseline that is independent of geographic location.

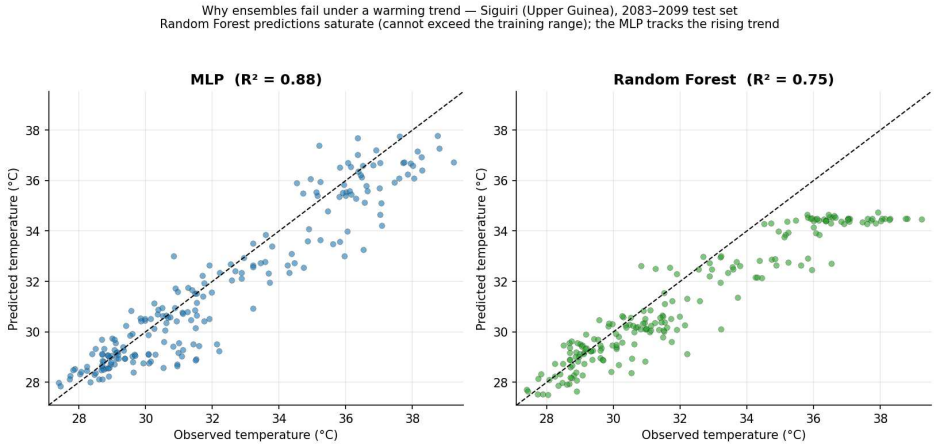


Fig. 3 Predicted versus observed monthly temperature on the 2083–2099 test set for Siguiri (Upper Guinea). The random forest (right) saturates at the upper edge of its training range and cannot reproduce the warmest months; the MLP (left) follows the warming trend. The dashed line is the 1:1 reference.

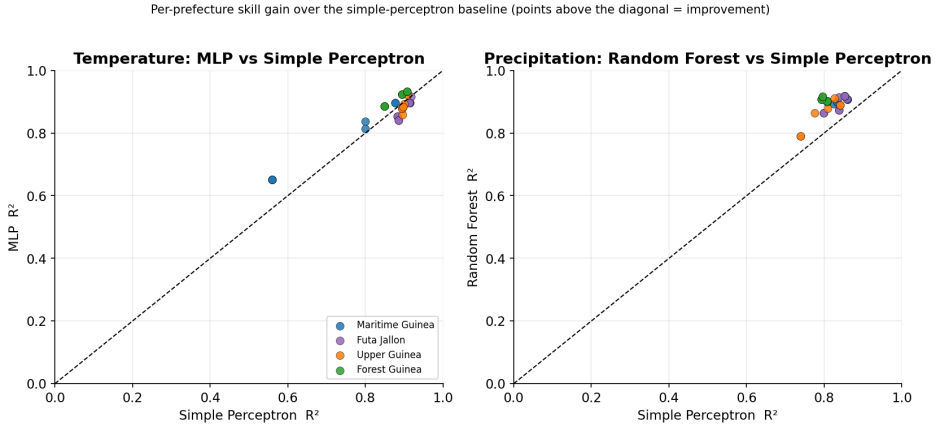


Fig. 4 Per-prefecture R^2 of the best nonlinear model against the simple-perceptron baseline (points above the diagonal indicate improvement). Left: temperature, MLP vs perceptron. Right: precipitation, random forest vs perceptron. Colours denote the four natural regions.

5.4 Regional patterns

The four natural regions impose systematic gradients on model skill (Table 2, Fig. 5). For temperature, the MLP's advantage over the perceptron is greatest in Maritime Guinea ($R^2 = 0.772$ versus 0.723) and in Forest Guinea ($R^2 = 0.910$ versus 0.879); in the more linearly behaved continental interior of Upper Guinea and the Fouta Djallon the linear baseline is already excellent and the neural network matches rather than exceeds it. The best single temperature result anywhere in the country is obtained

by the MLP in Guéckédou (Forest Guinea, $R^2 = 0.933$). The tree ensembles underperform in every region for temperature, most catastrophically along the coast. For precipitation the random forest is the strongest model in all four regions, peaking in Forest Guinea ($R^2 = 0.908$) and reaching its single best value in Labé (Fouta Djallon, $R^2 = 0.919$).

Table 2 Regional-mean R^2 (univariate, one-step-ahead). Neural/linear models lead for temperature; ensembles lead for precipitation, in every region.

Variable	Natural region	Simple Perceptron	MLP	Random Forest	Gradient Boosting
Temperature	Maritime Guinea	0.723	**0.772**	0.259	0.410
	Futa Jallon	**0.902**	0.879	0.736	0.769
	Upper Guinea	**0.898**	0.891	0.737	0.806
	Forest Guinea	0.879	**0.910**	0.637	0.704
Precipitation	Maritime Guinea	0.838	0.826	**0.901**	0.891
	Futa Jallon	0.838	0.793	**0.898**	0.886
	Upper Guinea	0.796	0.755	**0.865**	0.857
	Forest Guinea	0.800	0.801	**0.908**	0.903

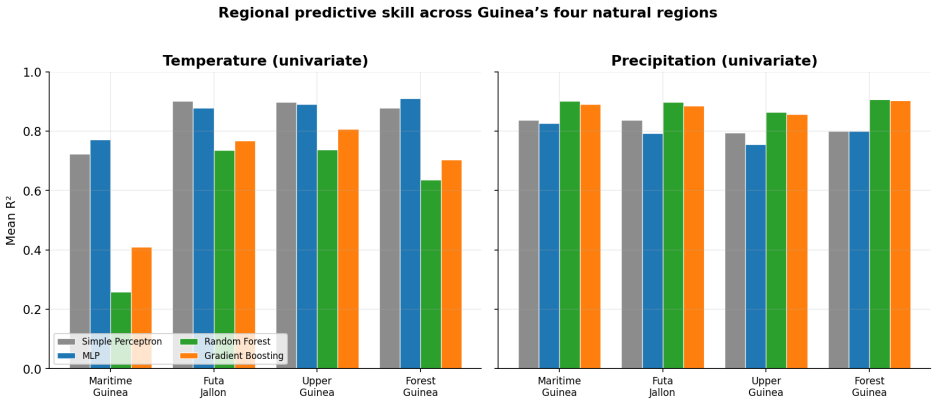


Fig. 5 Regional-mean R^2 by model for temperature (left) and precipitation (right) across Guinea's four natural regions.

5.5 Temporal behaviour of the forecasts

Figures 6 and 7 show observed and predicted series for six representative prefectures spanning the four regions, over the final years of the test period. For temperature (Fig. 6), the MLP reproduces both the seasonal cycle and the superimposed warming, with the closest agreement in the well-structured highland and forest prefectures (Mamou, Labé, Nzérékoré) and the widest residuals in coastal Conakry, whose near-constant maritime regime offers little exploitable seasonal variance. For precipitation

(Fig. 7), the random forest captures the timing and amplitude of the monsoon peak markedly better than the linear baseline, although, like all the models, it still tends to under-predict the most extreme wet-season totals, a residual limitation discussed in Section 8.

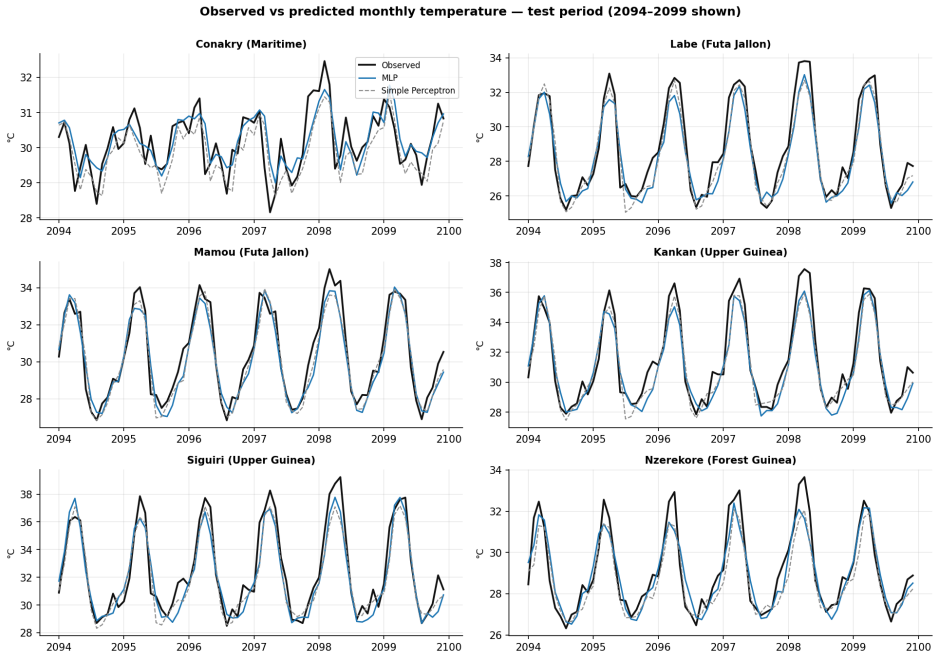


Fig. 6 Observed versus predicted monthly temperature (MLP and simple perceptron) for six representative prefectures, last six years of the test period.

5.6 Forecast horizon

At a 12-month lead time the advantage of the neural network over both competitors widens (Table 3, Fig. 8). Direct 12-month temperature forecasts retain useful skill for the MLP (national mean $R^2 = 0.759$), exceeding the simple perceptron (0.693) by about ten percent in relative terms, while the random forest degrades further to 0.450. The best 12-month results are again obtained in the Forest-Guinea prefectures (Guéckédou and Kissidougou, $R^2 \approx 0.91$), whose thermal stability and strong seasonal regularity make them especially predictable [59]. That the neural-network advantage is largest at the longer horizon, and in the forest region, is directly relevant to seasonal-to-annual forest planning [6].

5.7 Effect of the multivariate scenario

Adding the second variable's lags has small and asymmetric effects (Table 4). For temperature, the extra precipitation lags act mainly as noise: the MLP changes from

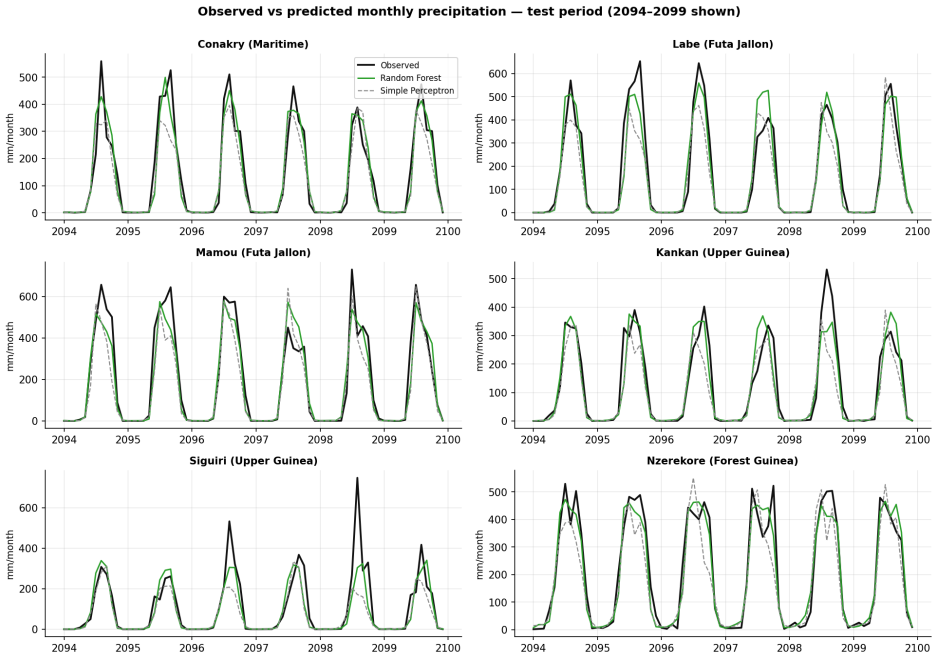


Fig. 7 Observed versus predicted monthly precipitation (random forest and simple perceptron) for six representative prefectures, last six years of the test period.

Table 3 Temperature forecast skill at 1- and 12-month horizons (national mean, direct multi-step forecasting). Values are R^2 with RMSE in °C in parentheses.

Horizon	**Simple Perceptron R^2 (RMSE)**	**MLP R^2 (RMSE)**	**Random Forest R^2 (RMSE)**
1 month	0.849 (0.79)	0.852 (0.80)	0.590 (1.29)
12 months	0.693 (1.11)	**0.759 (1.05)**	0.450 (1.42)

$R^2 = 0.86$ to 0.84 and the random forest degrades from 0.59 to 0.53. For precipitation, temperature lags carry weak positive information for the linear model (R^2 from 0.82 to 0.83) but do not improve the already-strong random forest. The univariate configuration is therefore preferred operationally for both variables, the random forest for precipitation and the MLP for temperature.

6 Discussion

The headline result is a clean split between the two families of nonlinear learners, and the statistical structure of the target variables accounts for it completely. Temperature in Guinea is a regular seasonal cycle riding on a strong, near-monotonic warming trend [1]. Both parts are smooth, and both are close to linearly predictable from twelve months of lags plus an explicit time index which is why the linear perceptron

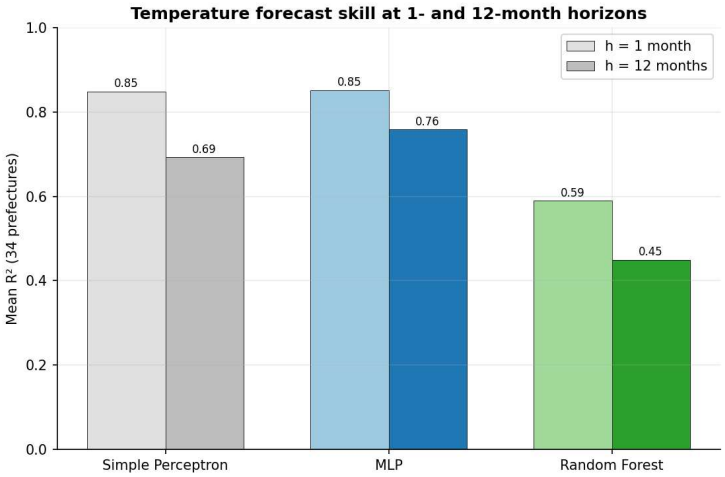


Fig. 8 Temperature forecast skill at 1- and 12-month horizons for the three principal models. The neural-network advantage widens at the longer lead time.

Table 4 Effect of the multivariate scenario on national-mean R^2 (one-step-ahead).

Variable	**Model**	**Univariate R^2**	**Multivariate R^2**	**Δ**
Temperature	Simple Perceptron	0.849	0.848	−0.001
	MLP	0.860	0.843	−0.017
	Random Forest	0.590	0.527	−0.063
	Gradient Boosting	0.669	0.665	−0.005
Precipitation	Simple Perceptron	0.820	0.830	+0.010
	MLP	0.794	0.758	−0.037
	Random Forest	0.893	0.882	−0.011
	Gradient Boosting	0.884	0.862	−0.022

already does well, and why the MLP, by mopping up the leftover nonlinear structure (the high-frequency coastal variability in particular), edges ahead to become the best model overall. But the factor that really decides the contest is extrapolation. The warming trend carries the test period past the upper edge of the training range, and only a model whose output is unbounded can keep climbing with it. Neural networks can; tree ensembles cannot [20], and their temperature skill collapses as a result.

Precipitation has the opposite character. It is bounded below by zero, highly skewed, dominated by a sharp seasonal switch between a near-dry and a strongly wet regime, and it carries no significant long-term trend [1]. This threshold-like, quasi-stationary, regime-switching structure is exactly what regression trees represent most naturally, through recursive binary splits, and the absence of a trend removes the extrapolation penalty [20]. The random forest therefore wins decisively for precipitation, while the smooth global function fitted by the MLP is comparatively ill-suited to the sharp regime boundaries and offers no advantage over the linear model.

These findings refine, rather than overturn, the conclusions of the earlier perceptron study [22]. That work correctly identified the linear architecture as the limiting factor and pointed to more expressive models; our results confirm that more expressive models help, but show that the right choice of expressive model is variable-specific, and that one popular and powerful family (tree ensembles) is actively counterproductive for trended variables. The widely held intuition that ensembles are a safe default for tabular prediction does not survive contact with a climate-change setting, where the most important predictions are, by definition, out-of-sample in the trend dimension [1, 24, 53].

The practical takeaway, then, is to stop looking for a single best model and instead match the tool to the variable: a multilayer perceptron for temperature, and for any variable that carries a secular trend; a random forest for precipitation, and for bounded, regime-switching variables that do not trend. A system built this way stays lightweight and fully reproducible, and it fits the data and the modest computing budgets that national meteorological and forestry services in West Africa actually have to work with.

7 Implications for forest-ecosystem monitoring

Guinea's forests (the montane and gallery forests of the Fouta Djallon, the dense humid forests of the south-east, and the woodland mosaics of the interior) are tightly coupled to the temperature and rainfall regimes analysed here [2–5]. Translating the forecasting results into monitoring practice yields several concrete implications.

First, temperature monitoring for forests should rely on neural networks. The variables most relevant to forest stress (the magnitude and timing of dry-season heat, the length of the warm season, and the secular rise in baseline temperature) are precisely those that the MLP captures and that tree ensembles miss because of their inability to extrapolate [20]. An ensemble-based temperature product would systematically under-predict future heat stress [1, 55], the most dangerous possible error for an early-warning system, whereas the neural network reproduces the warming faithfully, with its best skill achieved in the very Forest-Guinea prefectures (Guéckédou, Kissidougou, Nzérékoré) that contain the country's most valuable humid forests.

Second, precipitation and water-balance monitoring should rely on ensemble learning. The onset, peak and cessation of the monsoon determine the growing season, the fire-risk window and the recharge of the rivers that the Fouta Djallon supplies to much of West Africa [4]. The random forest's superior reproduction of monsoon timing and amplitude, uniform across all 34 prefectures, makes it the appropriate engine for rainfall-driven indicators such as season-onset alerts, dry-spell detection and simple water-balance accounting for forest hydrology [6].

Third, the longer-horizon results matter for planning. Because the neural-network advantage for temperature widens at the 12-month lead time, seasonal-to-annual outlooks for heat stress and phenological shifts (the timescales on which restoration, nursery and fire-management decisions are taken [6]) are best served by the MLP. The strong 12-month skill in the forest region ($R^2 \approx 0.91$) indicates that operationally useful annual temperature outlooks are achievable there with modest means.

Fourth, the regional structure of skill provides a map of confidence for monitoring agencies. Forecasts are most reliable in the highland and forest interior and least reliable along the oceanic coast, where high-frequency variability dominates [8]; monitoring products should carry this spatially varying uncertainty explicitly so that managers weight alerts accordingly. Taken together, these implications define a pragmatic, dual-engine climate-information layer that a national forest-monitoring system could adopt directly, using only historical series and commodity computing.

8 Limitations

Several limitations bound the interpretation of these results. The climate series are model-derived projections [7] rather than dense station observations; their internal smoothness makes them more predictable than raw gauge data and inflates absolute skill relative to an observational setting, although it does not affect the central, relative comparison among model classes. All models still under-predict the most extreme monthly rainfall totals, a well-known consequence of squared-error training on skewed targets that quantile or distributional loss functions could mitigate. The predictor set is deliberately restricted to lagged temperature and precipitation plus calendar features, omitting humidity, wind, radiation, soil moisture and large-scale teleconnection indices (for example the Atlantic Multidecadal Oscillation and West African Monsoon indices) whose inclusion would likely raise accuracy, particularly along the coast [8]. Each prefecture is modelled independently, ignoring the spatial correlation that convolutional or graph neural networks could exploit. Finally, richer architectures (recurrent networks such as LSTM and GRU [59], temporal convolutional networks, and hybrid physics-informed models [23, 60, 61]) were beyond the scope of this controlled four-model comparison and are natural avenues for future work, as is direct validation against remotely sensed indicators of forest condition [2].

9 Conclusion

We compared a simple perceptron, a multilayer perceptron and two ensemble-learning methods for prefecture-level monthly temperature and precipitation forecasting across all 34 prefectures of Guinea, using an 85-year record and a strict chronological protocol. The study establishes a clear, physically grounded division of labour between the leading families of nonlinear learners. Neural networks are the method of choice for temperature, where they achieve the highest and most spatially robust skill (national $R^2 = 0.86$, rising in relative advantage at the 12-month horizon) and, critically, where they alone can extrapolate the warming trend that causes tree ensembles to fail. Ensemble learning is the method of choice for precipitation, where the random forest reduces error by about a quarter relative to the linear baseline and reproduces the monsoon cycle uniformly across the country. The inability of tree ensembles to extrapolate emerges as a decisive and frequently overlooked criterion for any climate-change application. For forest-ecosystem monitoring in Guinea and the wider West African region, these findings recommend a lightweight, reproducible, dual-engine system (a neural network for temperature and heat stress, a random forest for rainfall and water

balance) that national forestry and meteorological services can deploy with the data and computing resources they already possess.

Supplementary information. Not applicable.

Acknowledgements. The authors thank the reviewers for their constructive comments.

Declarations

- **Funding:** The authors declare that no specific funding was received for this work.
- **Competing interests:** The authors declare no competing interests.
- **Ethics approval:** Not applicable.
- **Consent to participate:** Not applicable. This study does not involve any human participants.
- **Consent for publication:** Not applicable.
- **Data availability:** The monthly temperature and precipitation series for the 34 prefectures of Guinea (2015–2099) are available from the corresponding author on reasonable request.
- **Code availability:** The four learners were implemented from first principles in Python/NumPy; the code is available from the corresponding author on reasonable request.
- **Author contributions:** K.K.K. conceived and designed the study, developed the methodology, implemented the four learners in Python/NumPy, carried out the analysis, and drafted the manuscript. A.H.G.A. contributed the forest-ecosystem monitoring framework and the ecological interpretation of the results, and advised on the statistical and biometric aspects of the study design. M.N.C. contributed to data preparation and to the evaluation of the models. A.S. contributed to software validation and to the verification of the results. P.K. contributed to the literature review and to the interpretation of the regional findings. M.M. supervised the work, provided computing resources, and handled project administration. All authors reviewed, edited, and approved the final manuscript.

References

[1] IPCC. *Climate Change 2021: The Physical Science Basis*. Cambridge University Press, Cambridge, UK, 2021.

[2] M. C. Hansen, P. V. Potapov, R. Moore, et al. High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160):850–853, 2013.

[3] Y. Malhi, T. A. Gardner, G. R. Goldsmith, M. R. Silman, and P. Zelazowski. Tropical forests in the Anthropocene. *Annual Review of Environment and Resources*, 39:125–159, 2014.

[4] FAO. *Global Forest Resources Assessment 2020: Main report*. Food and Agriculture Organization of the United Nations, Rome, 2020.

[5] M. Brandt, C. J. Tucker, A. Kariryaa, et al. An unexpectedly large count of trees in the West African Sahara and Sahel. *Nature*, 587(7832):78–82, 2020. doi: 10.1038/s41586-020-2824-5.

[6] J. Sillmann et al. Understanding, modeling and predicting weather and climate extremes: Challenges and opportunities. *Weather and Climate Extremes*, 18: 65–74, 2017.

[7] V. Eyring, S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor. Overview of the CMIP6 experimental design and organization. *Geoscientific Model Development*, 9(5):1937–1958, 2016.

[8] F. Kaspar, A. Andersson, M. Ziese, and R. Hollmann. Contributions to the improvement of climate data availability and quality for Sub-Saharan Africa. *Frontiers in Climate*, 3:815043, 2022.

[9] H. Hersbach, B. Bell, P. Berrisford, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: 10.1002/qj.3803.

[10] S. H. Gebrechorkos, J. Leyland, L. Slater, et al. A high-resolution daily global dataset of statistically downscaled CMIP6 models for climate impact analyses. *Scientific Data*, 10:611, 2023. doi: 10.1038/s41597-023-02465-x.

[11] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.

[12] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[13] D. Maraun and M. Widmann. *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press, Cambridge, UK, 2018.

[14] W. W. Hsieh and B. Tang. Applying neural network models to prediction and data analysis in meteorology and oceanography. *Bulletin of the American Meteorological Society*, 79(9):1855–1870, 1998.

[15] H. R. Maier and G. C. Dandy. Neural networks for the prediction and forecasting of water resources variables. *Environmental Modelling & Software*, 15(1):101–124, 2000.

[16] C. Irrgang, N. Boers, M. Sonnewald, E. A. Barnes, C. Kadow, J. Staneva, and J. Saynisch-Wagner. Towards neural Earth system modelling by integrating artificial intelligence in Earth system science. *Nature Machine Intelligence*, 3(8): 667–674, 2021. doi: 10.1038/s42256-021-00374-3.

[17] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970):533–538, 2023. doi: 10.1038/s41586-023-06185-3.

[18] R. Lam, A. Sanchez-Gonzalez, M. Willson, et al. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023. doi: 10.1126/science.adi2336.

[19] P. Bauer, P. D. Dueben, T. Hoefler, T. Quintino, T. C. Schulthess, and N. P. Wedi. The digital revolution of Earth-system science. *Nature Computational Science*, 1(2):104–113, 2021. doi: 10.1038/s43588-021-00023-0.

[20] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[21] J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232, 2001.

[22] K. K. Katakpe, M. N. Camara, A. Sow, P. Kpoghomou, and M. Makanera. Modeling and predicting climate parameters in Guinea using simple perceptrons: A study on temperature and precipitation. *International Journal of Computer Applications*, 187(101):41–58, 2026.

[23] T. Schneider, L. R. Leung, and R. C. Wills. Opinion: Optimizing climate models with process knowledge, resolution, and artificial intelligence. *Atmospheric Chemistry and Physics*, 24(12):7041–7062, 2024.

[24] S. Materia et al. Artificial intelligence for climate prediction of extremes: State of the art, challenges, and future perspectives. *WIREs Climate Change*, 15(6): e914, 2024.

[25] D. Rolnick, P. L. Donti, L. H. Kaack, et al. Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2):1–96, 2022. doi: 10.1145/3485128.

[26] J. Willard, X. Jia, S. Xu, M. Steinbach, and V. Kumar. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4):1–37, 2022. doi: 10.1145/3514228.

[27] G. A. Meehl, J. H. Richter, H. Teng, et al. Initialized Earth system prediction from subseasonal to decadal timescales. *Nature Reviews Earth & Environment*, 2(5):340–357, 2021. doi: 10.1038/s43017-021-00155-x.

[28] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

[29] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[30] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, 1969.

[31] K.-I. Funahashi. On the approximate realization of continuous mappings by neural networks. *Neural Networks*, 2(3):183–192, 1989.

[32] K. Hsu, H. V. Gupta, and S. Sorooshian. Artificial neural network modeling of the rainfall–runoff process. *Water Resources Research*, 31(10):2517–2530, 1995.

[33] S. Haykin. *Neural Networks and Learning Machines*. Pearson, New York, 3rd edition, 2009.

[34] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[35] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, Cambridge, MA, 2016.

[36] T. Nguyen, J. Brandstetter, A. Kapoor, J. K. Gupta, and A. Grover. ClimaX: A foundation model for weather and climate, 2023. International Conference on Machine Learning (ICML). arXiv:2301.10343.

[37] T. Beucler, M. Pritchard, S. Rasp, J. Ott, P. Baldi, and P. Gentine. Enforcing analytic constraints in neural networks emulating physical systems. *Physical Review Letters*, 126(9):098302, 2021. doi: 10.1103/PhysRevLett.126.098302.

[38] K. Kashinath, M. Mustafa, A. Albert, et al. Physics-informed machine learning: case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A*, 379(2194):20200093, 2021. doi: 10.1098/rsta.2020.0093.

[39] S.-Q. Dotse et al. Deep learning-based long short-term memory recurrent neural networks for monthly rainfall forecasting in Ghana, West Africa. *Theoretical and Applied Climatology*, 155(4):3033–3045, 2024.

[40] N. K. A. Appiah-Badu, Y. M. Missah, L. K. Amekudzi, et al. Rainfall prediction using machine learning algorithms for the various ecological zones of Ghana. *IEEE Access*, 10:5069–5082, 2021.

[41] A. T. Ogunrinde, P. G. Oguntunde, J. T. Fasinmirin, and A. S. Akinwumiju. Application of artificial neural network for forecasting standardized precipitation and evapotranspiration index: A case study of Nigeria. *Engineering Reports*, 2 (7):e12194, 2020.

[42] A. Aizansi, K. Ogunjobi, and F. Ogou. Monthly rainfall prediction using artificial neural network (case study: Republic of Benin). *Environmental Data Science*, 3: e10, 2024.

[43] H. Citakoglu. Comparison of multiple learning artificial intelligence models for estimation of long-term monthly temperatures in Turkey. *Arabian Journal of Geosciences*, 14(20):2131, 2021.

[44] Q. Guo, Z. He, and Z. Wang. Monthly climate prediction using deep convolutional neural network and long short-term memory. *Scientific Reports*, 14(1):17748, 2024.

[45] Y. Fan, V. Krasnopolsky, H. Van Den Dool, C.-Y. Wu, and J. Gottschalck. Using artificial neural networks to improve CFS week-3–4 precipitation and 2-m air temperature forecasts. *Weather and Forecasting*, 38(5):637–654, 2023.

[46] T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

[47] I. Maqsood et al. Statistical downscaling of temperature and precipitation using machine learning under multiple climate scenarios. *International Journal of Climatology*, 2022. Comparative downscaling study.

[48] R. Balmaceda-Huarte et al. Convolutional neural networks for the downscaling of daily temperatures over South America. *International Journal of Climatology*, 2024.

[49] Omkar S. Kalange, Rushikesh S. Kahat, Atharv S. Kale, Trupti R. Kale, and Pushkar S. Joglekar. Implementation of various machine learning algorithms for traffic sign detection and recognition. In *Proceedings of the International Conference on Computing, Communication and Networking Technologies*, 2022.

[50] Z. Shen et al. Current progress in subseasonal-to-decadal prediction based on machine learning. *Applied Computing and Geosciences*, 24:100201, 2024.

[51] D. Watson-Parris, Y. Rao, D. Olivié, et al. ClimateBench v1.0: A benchmark for data-driven climate projections. *Journal of Advances in Modeling Earth Systems*, 14(10):e2021MS002954, 2022. doi: 10.1029/2021MS002954.

[52] E. A. Barnes, B. Toms, J. W. Hurrell, I. Ebert-Uphoff, C. Anderson, and D. Anderson. Indicator patterns of forced change learned by an artificial neural network. *Journal of Advances in Modeling Earth Systems*, 12(9):e2020MS002195, 2020. doi: 10.1029/2020MS002195.

[53] V. Eyring, W. D. Collins, P. Gentine, et al. Pushing the frontiers in climate modelling and analysis with machine learning. *Nature Climate Change*, 14:916–928, 2024. doi: 10.1038/s41558-024-02095-y.

[54] A. Dosio, M. W. Jury, M. Almazroui, et al. Projected future daily characteristics of African precipitation based on global (CMIP5, CMIP6) and regional

(CORDEX, CORDEX-CORE) climate models. *Climate Dynamics*, 57:3135–3158, 2021. doi: 10.1007/s00382-021-05859-w.

[55] M. B. Sylla, A. Faye, F. Giorgi, A. Diédhióu, and H. Kunstmann. Projected heat stress under 1.5°C and 2°C global warming scenarios creates unprecedented discomfort for humans and key agroeconomic sectors in West Africa. *Earth's Future*, 10(1):e2021EF002234, 2022. doi: 10.1029/2021EF002234.

[56] A. A. Akinsanola, K. O. Ogunjobi, V. O. Ajayi, et al. Evaluation of rainfall simulations over West Africa in dynamically downscaled CMIP5 models. *Climate Dynamics*, 55:2253–2271, 2020. doi: 10.1007/s00382-020-05395-3.

[57] S. Lawal, C. Lennard, C. Jack, P. Wolski, B. Hewitson, and B. Abiodun. Evaluating CMIP6 model simulations of mean precipitation and temperature over the West African monsoon region. *International Journal of Climatology*, 43: 4199–4218, 2023. doi: 10.1002/joc.8103.

[58] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. arXiv:1412.6980.

[59] B. Lim, S. O. Arik, N. Loeff, and T. Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, 2021.

[60] S. Rasp, M. S. Pritchard, and P. Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018.

[61] T. Weber, A. Corotan, B. Hutchinson, B. Kravitz, and R. Link. Deep learning for creating surrogate models of precipitation in Earth system models. *Atmospheric Chemistry and Physics*, 20(4):2303–2317, 2020.

Appendix A Prefecture-level results

for all four models in the univariate, one-step-ahead scenario, together with the simple-perceptron RMSE. Rows are ordered by natural region. The national mean is shown in the final row. SP = simple perceptron; MLP = multilayer perceptron; RF = random forest; GB = gradient boosting.

Table A1 Monthly temperature prediction (°C) per-prefecture R^2 and simple-perceptron RMSE.

Region	Prefecture	SP RMSE (°C)	SP R^2	MLP R^2	RF R^2	GB R^2
Maritime Guinea	Boké	0.70	0.799	0.837	0.507	0.673
	Boffa	0.63	0.560	0.652	−0.234	0.026
	Fria	0.63	0.560	0.652	−0.234	0.026
	Kindia	0.61	0.877	0.897	0.724	0.751
	Dubreka	0.63	0.560	0.652	−0.234	0.026
	Coyah	0.61	0.877	0.897	0.724	0.751
	Forécariah	0.54	0.800	0.813	0.544	0.600
	Conakry	0.63	0.560	0.652	−0.234	0.026
	Télimélé	0.67	0.915	0.896	0.770	0.814
Futa Jallon	Gaoual	0.81	0.883	0.853	0.689	0.741
	Koundara	0.81	0.883	0.853	0.689	0.741
	Labé	0.80	0.913	0.896	0.741	0.797
	Lélouma	0.67	0.915	0.896	0.770	0.814
	Tougué	0.80	0.913	0.896	0.741	0.797
	Mali	0.93	0.886	0.841	0.732	0.712
	Dalaba	0.80	0.913	0.896	0.741	0.797
	Pita	0.80	0.913	0.896	0.741	0.797
	Mamou	0.69	0.917	0.917	0.785	0.777
	Koubia	0.93	0.886	0.841	0.732	0.712
Upper Guinea	Dinguiraye	0.95	0.897	0.885	0.768	0.820
	Kouroussa	1.00	0.896	0.859	0.763	0.799
	Siguiri	1.04	0.894	0.879	0.745	0.833
	Kankan	0.94	0.901	0.894	0.709	0.818
	Mandiana	1.04	0.894	0.879	0.745	0.833
	Faranah	0.80	0.909	0.923	0.718	0.780
	Kérouané	0.85	0.894	0.924	0.684	0.749
	Dabola	0.95	0.897	0.885	0.768	0.820
Forest Guinea	Beyla	0.85	0.894	0.924	0.684	0.749
	Lola	0.77	0.850	0.886	0.523	0.636
	Nzérékoré	0.77	0.850	0.886	0.523	0.636
	Yomou	0.77	0.850	0.886	0.523	0.636
	Macenta	0.85	0.894	0.924	0.684	0.749
	Guéckédou	0.79	0.907	**0.933**	0.762	0.761
	Kissidougou	0.79	0.907	0.933	0.762	0.761
National mean	34 prefectures	0.79	0.849	**0.860**	0.590	0.669

Table A2 Monthly precipitation prediction (mm) per-prefecture R^2 and simple-perceptron RMSE.

Region	Prefecture	SP RMSE (mm)	SP R^2	MLP R^2	RF R^2	GB R^2
Maritime Guinea	Boké	77.60	0.825	0.844	0.893	0.863
	Boffa	69.85	0.834	0.825	0.899	0.894
	Fria	69.85	0.834	0.825	0.899	0.894
	Kindia	80.27	0.860	0.836	0.907	0.907
	Dubreka	69.85	0.834	0.825	0.899	0.894
	Coyah	80.27	0.860	0.836	0.907	0.907
	Forécariah	72.38	0.807	0.791	0.899	0.878
	Conakry	69.85	0.834	0.825	0.899	0.894
	Télimélé	94.83	0.858	0.830	0.910	0.882
Futa Jallon	Gaoual	74.44	0.799	0.812	0.865	0.855
	Koundara	74.44	0.799	0.812	0.865	0.855
	Labé	81.77	0.852	0.762	**0.919**	0.910
	Lélouma	94.83	0.858	0.830	0.910	0.882
	Tougué	81.77	0.852	0.762	0.919	0.910
	Mali	62.08	0.838	0.813	0.873	0.866
	Dalaba	81.77	0.852	0.762	0.919	0.910
	Pita	81.77	0.852	0.762	0.919	0.910
	Mamou	90.81	0.839	0.795	0.915	0.900
	Koubia	62.08	0.838	0.813	0.873	0.866
Upper Guinea	Dinguiraye	57.69	0.842	0.843	0.889	0.890
	Kouroussa	64.82	0.810	0.778	0.878	0.866
	Siguiri	77.29	0.739	0.640	0.790	0.780
	Kankan	76.25	0.776	0.695	0.865	0.845
	Mandiana	77.29	0.739	0.640	0.790	0.780
	Faranah	79.18	0.826	0.801	0.911	0.899
	Kérouané	93.79	0.792	0.800	0.908	0.903
	Dabola	57.69	0.842	0.843	0.889	0.890
Forest Guinea	Beyla	93.79	0.792	0.800	0.908	0.903
	Lola	86.01	0.808	0.817	0.902	0.901
	Nzérékoré	86.01	0.808	0.817	0.902	0.901
	Yomou	86.01	0.808	0.817	0.902	0.901
	Macenta	93.79	0.792	0.800	0.908	0.903
	Guéckédou	106.23	0.797	0.778	0.916	0.906
	Kissidougou	106.23	0.797	0.778	0.916	0.906
National mean	34 prefectures	79.78	0.820	0.794	**0.893**	0.884