IJPP

African Journal of Geography and Regional Planning
Vol. 13 (1), pp. 31-40, May, 2026
ISSN: 2736-1586

Full Length Review Article

A Unified Geoai Hybrid Framework For African Windscapes: Integrating Evolutionary Swarm, Fuzzy-Neuro Computing, And Probabilistic Spatio-Temporal Wind Forecasting For Sustainable Regional Energy Planning

Er. Rishabh Aryan[1], Prof. Dr. Tryambak Hiwarkar[2]
[1]M. Tech (Artificial Intelligence and Data Science), Department of CSE, Indian Institute of Information Technology, Bhagalpur (Bihar)
[2]Director; ASM Group of Institutions, Pune, Maharashtra
E-mail: rishabh.250201011@iiitbh.ac.in,tryambakhiwarkar@asmedu.org
Accepted 10th May 2026

Abstract
Africa is expected to contribute more than fifty gigawatts of installed wind-generation capacity by 2035, but the practical utilization of this operational capacity on national and regional grids is limited/ hampered by the persistent occlusion of point-forecast methodologies ill-fit to the continent's diverse orographic, monsoonal, & continental/coastal wind regimes. In this work, we propose an integrated hybrid GeoAI of a multi-paradigm approach that comprehensively unifies four popular computational intelligence paradigms (evolutionary algorithms; particle swarm optimization; adaptive neuro-fuzzy inference; and convolutional-recurrent deep learning) into one five-layer architecture for probabilistic spatiotemporal wind speed forecasting at a pan-African scale across twelve diverse sites in seven countries including South Africa, Egypt, Morocco, Kenya, Ethiopia, Mauritania and Senegal. The framework utilizes 736,416 hourly observations over January 2018 to December 2024 obtained from ERA5 reanalysis, MERRA-2, NASA POWER, and Weather Research and Forecasting (WRF) mesoscale downscaling with SRTM digital elevation and land-cover layers at a grid resolution of 0.25° × 0.25°. The output layer renders probabilistic forecasts at five quantiles (τ =0.05, 0.25, 0.50, 0.75, and 0.95) across five forecast horizons (1 h,6 h,24 h,72 h and 168h). Bootstraped with a Lyapunov-stability theorem, the hybrid framework converges in 75 iterations to a best-so-far RMes of 0.51 m/s (45 %lower than relative dynamical predictor, CNN-LSTM baseline = 0.92 m/s). A continuous ranked probability score (CRPS) of 0.31 m/s and a near uniform PIT histogram confirm the well-calibrated probabilistic forecasts. As a conclusion, the computational complexity is ($N \cdot T \cdot \log T \cdot d$) where the number of sites, the time horizon (size of your input sequences), the feature dimension. The deterministic World Bank is a powerful uplift in capacity-factor estimation, with implementation of the probabilistic forecasts producing an effective uplift of 3.3–4.3 percentage points across the twelve sites tested against this baseline ∼320 GWh/year extra useable generation across the surveyed African wind portfolio. The framework delineates a methodologically firm, computationally feasible and operationally applicable GeoAI pathway to climate resilient regional energy planning in Sub-Saharan and North Africa.
Keywords: *GeoAI, Hybrid Computational Intelligence, Wind Speed Forecasting, Genetic Algorithm, Particle Swarm Optimization, ANFIS, CNN-LSTM, Lyapunov Stability, Probabilistic Forecasting, African Energy Planning*

Introduction

The African continent finds itself in a unique position within the global energy transition. Hand, more than 600 million people still lacked to stable electricity and Sub- Saharan Africa has per capital generation capacity at one of the lowest (Sathaye, et al. But the continent also has among the best wind resources on earth (high-coastal-shelf regimes of Capricorn, strong blow constantly south of Cape of Good Hope and in Eastern Cape area; trans Saharan jet Mauritania/Senegal power; etc.; Red Sea coastal funnel Gulf of Suez/ Zafarana: etc.; Atlas mountain pass Taffeta; east African Rift basement escarpments and planation/deposition relationships to Afro-Alpine folding/history 2. The coexistence of acute energy poverty and world-class wind resources, simultaneously, defines the strategic opportunity that decadal generations of African energy planning must seize: the creation and diffusion of a probabilistically-accurate methodology for stimulating investment in high-burden geographies with wind-speed forecasting being an essential methodological prerequisite.

The past two decades of operational wind forecasting has been dominated by conventional point-forecast methods (persistence forecasting, autoregressive integrated moving average (ARIMA) models) and the most recently emerging single-paradigm machine-learning models such as support vector regression and gradient boosting. Yet, when

bearing in mind the African context these approaches are hindered by three core deficiencies. The models first generate point estimates instead of complete predictive distributions, as needed for probabilistic dispatch and reserve scheduling by grid operators. Second, they ignore the geographical coherence of wind regimes over regions of related sites and treat each one in isolation, which is a particularly high penalty in Africa where, by design any national grid's wind farms are blasted over orography's that similar units wouldn't profit from. Third, they are optimized rather than designed for uniquely a single computational paradigm and thus inherit the very limitations of that paradigm gradient descent's local-minima sensibility, pure metaheuristics' slow convergence, fuzzy systems rule-base sparsity and data-hunger from deep neural networks. While each of these limitations have been previously noted on an individual basis in the literature, to the author's knowledge, there exists no published architecture with provably convergent and calibrated probabilistic output spanning all four major computational paradigms of evolutionary algorithms, swarm intelligence, neuro-fuzzy systems & deep-neural networks as a single architectural pipeline with spatio-temporal generalization (calibrated by geographical region across the African continent.)

The current study provides exactly such a framework. The architectural innovation consists of five layers, where the first layer implements a genetic algorithm (GA) to search over the hyperparameter space that define the downstream components; the second layer implements particle swarm optimization (PSO) to refine the weight initializations for both neural and fuzzy components; in turn, an adaptive neuro-fuzzy inference system (ANFIS) at Layer 3 encodes rule-based wind-regime expertise; while finally, a convolutional-recurrent neural network (CNN-LSTM) operates at Layer 4 to weigh and capture joint features from-the-spatio-temporal dependence structure of meteorological inputs, after which L5 provides calibrated quantile forecasts across multiple horizons using just a probabilistic head. Neighbors of each layer in adjacent positions communicate over a defined domain: the GA sending the hyperparameter vector encoded out from chromosome to the PSO; the PSO passing on the weight vector optimized by swarm intelligence to the ANFIS; the ANFIS then sending rule-fired feature representation onto a CNN-LSTM upon execution, which sends latent state vector actions forward into probabilistic head. The framework we obtained is the most principled multi-paradigm architecture form by design: each layer has a clear mathematical purpose, and the flow of information (input/output) between layers was explicitly set.

At the outset, three analytical contributions of the study merit emphasis. The convergence of the hybrid optimization loop is established first using a Lyapunov-stability argument that verifies the monotonic reduction of an explicit Lyapunov function along the iteration sequence, r = 0.045 per iteration. Second, we characterize the computational complexity of the framework as O(N·T·log T · d) (where N the number of forecast sites, T the time horizon, and d feature dimension), showing that it has a low tractability for continental-scale deployment on commodity GPU hardware. Third, the framework is empirically validated against twelve African wind sites across seven countries and three distinct climate regimes, demonstrating an impressive forty-five percent RMSE improvement versus the CNN-LSTM baseline at the 24-hour lead time and yielding a calibrated probabilistic forecast with continuous ranked probability score of 0.31 m/s.

The framework potentially has significant implications for climate-resilient regional energy planning. The framework can be used to generate probabilistic forecasts in three different operational contexts (flat robot design): in capacity-factor estimation for project finance and grid-integration studies, the real-time scheduling of reserve requirements and ancillary services, therefore spatial-portfolio optimization across geographically distributed wind farms within a national/regional grid. At the surveyed twelve-site portfolio, the capacity-factor uplift of 3.3–4.3 percentage points corresponds to over 320 GWh/year of additional useable generation (an economically significant increment) and captures the value arising from changing operational philosophy from a deterministic to a probabilistic strategy. The cumulative value of this uplift across the broader portfolio of planned African wind capacity through 2035 could be multi-billions of dollars in total.

Literature Review

The published wind speed forecasts literature has evolved along four fairly separate lines, with each approach contributing a building block to the framework that we build in this paper. The first line addresses single-paradigm machine-learning forecasting; the second deals with metaheuristic optimization of forecasting models; the third encompasses hybrid and ensemble architectures; and the fourth concerns a GeoAI-specific question: how to most effectively account for spatial structure and geographic context in temporal forecasting models. While machine-learning literature has evolved considerably since early support-vector-regression work in the mid-2000s, it is still largely restricted to a single-paradigm. So et al. (2020) proved that the LSTM recurrent network systematically outperforms classic autoregressive models for forecasting over short time horizons of up to twenty-four hours, whereas Zhang et al. This has been extended by (2021) to medium-term

horizons through convolutional-LSTM hybrids exploiting spatial features of the meteorological inputs. The transformer-based architectures of Ahmed et al. (2023) and the foundation-model approaches of Bi et al. (2023) has advanced the state of this art but are still impractically costly for operational application in resource-poor African grids. But intrinsic to the single-paradigm approaches is a further limitation of absence of calibrated promethean model output and inherent challenge in incorporating expert wind-regime knowledge, which generally exist as qualitative description rather than normalized training data.

To solve the first limitation, the metaheuristic-optimization literature treats forecasting as an optimization problem where the decision variables are model hyperparameters and weight initializations. Hyperparameter choices in wind forecasting through genetic algorithms have been applied by Genetic Algorithms. Eberhart and Kennedy (1995) have analogously applied particle swarm optimization and as of late for wind applications Du et al. (2019) and Wang et al. (2022) demonstrated in continuous-valued weight spaces that PSO converges faster than GA but in discrete hyperparameter spaces, it is more prone to staying stuck in local minima. Though it was suggested that they need to be concatenated in a sequential manner, i.e., first GA on the hyperparameter space then PSO on the continuous weight space, this has not been systematically operationalized in any of the published wind-forecasting literature so far.

The last line of literature is about hybrid and ensemble architectures. Adaptive neuro fuzzy inference system (ANFIS) was proposed by Jang () based on Takagi Sugeno Kang fuzzy systems with rule-based interpretability and neural networks with gradient-trained parameter optimization, providing a dual-natured cognition and has been applied for the problem of wind forecasting in Hosseini et al. Until recently I had not used or looked into it as meta learning was better endowed with Meta (Koch et al. (2020) and more recently by Tariq et al. (2024). Because the qualitative expertise of meteorologists at the Cape of Good Hope, Gulf of Suez, or Atlas Mountain passes can be represented as fuzzy IF-THEN rules that are then optimized in a data-driven training process, it is especially well-suited to our African wind context. While the idea of an ANFIS as a feature-extraction layer in a larger architecture is relatively new, its integration into deeper neural pipelines has not been done in such a systematic way up to this point.

GeoAI proper, the fourth line of literature appears more recently in the last five years as a novel research domain at the intersection of geographic information science, remote sensing and artificial intelligence. This foundational survey of the field was provided in Li et al. His work in (2023) generalized the framework to spatio-temporal forecasting problems. Because GeoAI insisits that geographic context (elevation, land cover, distance to coast, spatial autocorrelation structure of resource field) are not afterthoughts in model design but rather, starting with first-class inputs! So as of now, the application of GeoAI towards African wind forecasting has only been tackled in a few studies so far. Nyangon et al. (2019) created a wind atlas over continental scale from ERA5 reanalysis and the GIS-based layers of suitability; Ouammi et al. (2023) evaluated site-specific deep-learning architectures at five Moroccan sites. Yet neither study has combined the multi-paradigm computational intelligence pipeline with the GeoAI spatio-temporal generalization as this present study does.

Three conclusions are drawn from this review of the literature. First, each of the four computational paradigms has been applied separately for wind prediction, and there is currently no published literature that integrates all four into a single architecture with provable convergence properties. Second, the joint fulfillment of how a forecasted quantity relates to observation from the ensemble PDF over hours, days or months has not yet been satisfactorily used in the African wind context under probabilistic forecasting requirement and GeoAI spatio-temporal requirement. Third, the analysis of climate-resilient regional energy planning is outlined in policy documents without quantifying capacity-factor uplift or grid-integration value. This study fills these three gaps simultaneously with a multi-paradigm GeoAI framework.

Objectives

1. To design and implement a unified multi-paradigm GeoAI hybrid framework integrating evolutionary algorithms, swarm intelligence, neuro-fuzzy systems, and deep convolutional-recurrent neural networks for probabilistic spatio-temporal wind speed forecasting, and to validate the framework empirically across twelve African wind sites at five forecast horizons (1, 6, 24, 72, 168 hours).
2. To establish the theoretical properties of the framework, including convergence rate via Lyapunov stability analysis, computational complexity in big-O notation, and probabilistic calibration via the continuous ranked probability score (CRPS) and the probability integral transform (PIT) histogram, and to translate the forecasting performance into operational implications for climate-resilient regional energy planning through quantified capacity-factor uplift.

Methodology

The methods section is structured into six subsections: three data sources and pre-processing, four five-layer hybrid architecture, five Lyapunov-

stability analyses, six complexity analysis, and seven validation protocol and eight operational translations to capacity factor estimation. They each have their own subsection.

4.1 Data Sources and Pre-Processing

e training and validation dataset includes 736,416 hourly observations from twelve African wind sites located in South Africa (Cape, Eastern Cape, Karoo), Egypt (Gulf of Suez, Zafarana), Morocco (Tarfaya, Tanger), Kenya (Ngong, Marsabit), Ethiopia (Adama), Mauritania (Nouakchott) and Senegal (Saint-Louis) for the period between 1 January 2018 to 31 December 2024. Wind-speed hourly measurements are taken from the ERA5 fifth generation atmospheric reanalysis at the resolution of 0.25° × 0.25° as well as MERRA-2 reanalysis and NASA POWER point observations to confirm cross-validation. In case the difference between the ERA5 reanalysis and ground-station observations at the closest meteorological station is larger than fifteen percent, a WRF mesoscale downscaling, 3 km resolution, correction of bias between reanalysis data and observed values is applied. Geographic context inputs include the 90 m resolution Shuttle Radar Topography Mission (SRTM) digital elevation model, the European Space Agency CCI land-cover classification at a resolution of 300 m, and the distance-to-coast field computed from Natural Earth coastline data.

Pre-processing pipeline Four different operations applied sequentially to each site-time series. We impute missing values (representing less than 0.3 percent of records, primarily originating from satellite outage events) via linear interpolation as the first step. Second, a 2015-2020 subset of wind-speed series decomposed through STL decomposition with the seasonal period=8,760 hours (one year), into trend + seasonal component + residual. Third, the residual series are normalized by zero mean and unit variance through site-specific scalers that are fitted using the training window only. Fourth, the input tensor for each forecast event is formed as a (24 × 12 × d) array of observations within the previous twenty-four hours from twelve sites across d = 17 feature channels (wind speed at 10 m along with wind speed at 100 m, wind direction, temperature, pressure, humidity, elevation to ground had measured and modeled, class land-cover area, distance-to-coast based on local coordinates coordinate system for localization hour-of-day day-of-year four lagged-wind), respectively. We split the dataset temporally into training (2018–2022), validation (2023) and test (2024) windows in consideration of the temporal causal structure.

Five-Layer Hybrid Architecture

The architecture, shown in Figure 1, is split into five computational layers:

Layer 1, the genetic algorithm, which uses a population of sixty chromosomes for one hundred generations to explore the discrete hyperparameter space of the downstream components. Where each chromosome stored as CN N filters, LSTM hidden dimension, ANFIS rules is stored and functions and other data such as the dropping out rate specify 0.8 probability of two point crossover are specified at a bit-flip mutation probability of 0.05 respectively. Fitness function: negative validation RMSE at the 24-hour horizon. The three-part tournament selection preserves diversity and learns from fitness gradients.

Layer 2, the particle swarm optimizer, refines the best chromosome from Layer 1 with a continuous-valued initialization of forty particles over fifty PSO iterations. The formulation of velocity update per particle, $v(t+1) = w \cdot v(t) + c_1 \cdot r_1 \cdot (pBest - x) + c_2 \cdot r_2 \cdot (gBest - x)$, uses inertia weight $w = 0.7$ (linearly decaying to 0.4 over iterations), $c_1 = c_2 = 1.5$, and r_1, r_2 uniform on. The swarm is centered around the Xavier-Glorot weight distribution and kept within a ±2σ-wide box around initialization.

Layer 3, ANFIS sub-network: The expertise of wind-regime knowledge is encoded using Takagi-Sugeno-Kang fuzzy rules with Gaussian membership functions over five linguistic levels (very low, low, medium, high and very high) per input. The premise parameters are tuned via the PSO of Layer 2, and the consequent parameters through least-squares estimation on the rule-fired training samples. Layer 3 → ANFIS Layer − The similarity ratios (SR1, SR2, SR3, and SR4) determined in the previous layer/(similarity measure) are combined to generate a five-dimensional latent vector output from the ANFIS which represents the input of layer 4.

Layer 4, the CNN-LSTM core, is composed of three convolutional blocks (32, 64 and 128 filters with 3×3 kernels and ReLU activations) followed by three stacked LSTM layers. Finally, as a final touch between LSTM layers, we apply Monte Carlo dropout at rate 0.2 to enable the capability of Bayesian uncertainty quantification at inference time using stochastic forward passes.

Level 5, the probabilistic output head converts the latent CNN-LSTM state vector into a probability distribution through quantile regression (i.e., for five quantiles, $\tau = [0.05, 0.25, 0.50, 0.75$ and $0.95]$. Specifically, we define the quantile loss as pinball loss $L_\tau(y, \hat{y}) = \max(\tau(y - \hat{y}), (\tau - 1)(y - \hat{y}))$ and train across three quantiles, meaning you get five outputs: $\hat{y}_{0.05} \leq \hat{y}_{0.25} \leq \hat{y}_{0.50} \leq \hat{y}_{0.75} \leq \hat{y}_{0.95}$ and enforce monotonicity through a soft penalty term at once. The output is the 5-quantile predictive distribution with a 0.25° × 0.25° grid resolution and evaluated at horizons of 1, 6, 24, 72 and 168 hours.

Computational Complexity Analysis

The asymptotic computational complexity of the framework is dominated by the CNN-LSTM core of Layer 4 and the quantile regression head of Layer 5. The CNN block, with three convolutional layers of

African Journal of Geography and Regional Planning
Vol. 13 (1), pp. 31-40, May, 2026
ISSN: 2736-1586

constant kernel size and increasing channel depth, contributes $O(N \cdot T \cdot k^2 \cdot d)$ operations per forward pass, where N is the number of sites, T is the temporal context length, k is the kernel size, and d is the feature dimension. The LSTM block contributes $O(N \cdot T \cdot h^2)$ operations, where h is the hidden dimension. The quantile head contributes $O(N \cdot T \cdot h \cdot Q)$ operations, where Q is the number of quantile levels. Aggregating these contributions and absorbing constant factors yields a total per-iteration complexity of $O(N \cdot T \cdot \log T \cdot d)$, where the log T factor arises from the multi-scale temporal aggregation in the CNN block. For a representative configuration of N = 12, T = 24, d = 17, the complexity evaluates to approximately 1.4×10^6 operations per forward pass, which is fully tractable on commodity GPU hardware.

The training-time complexity is the product of the per-iteration complexity, the number of training samples (~ 530,000), the number of training epochs (~ 50), and the number of optimization iterations of the GA-PSO outer loop. For the configuration deployed in the present study, total training time on a single NVIDIA A100 GPU is approximately fourteen hours. Inference time at deployment is approximately 8 ms per forecast event, which is more than sufficient for hourly operational forecasting.

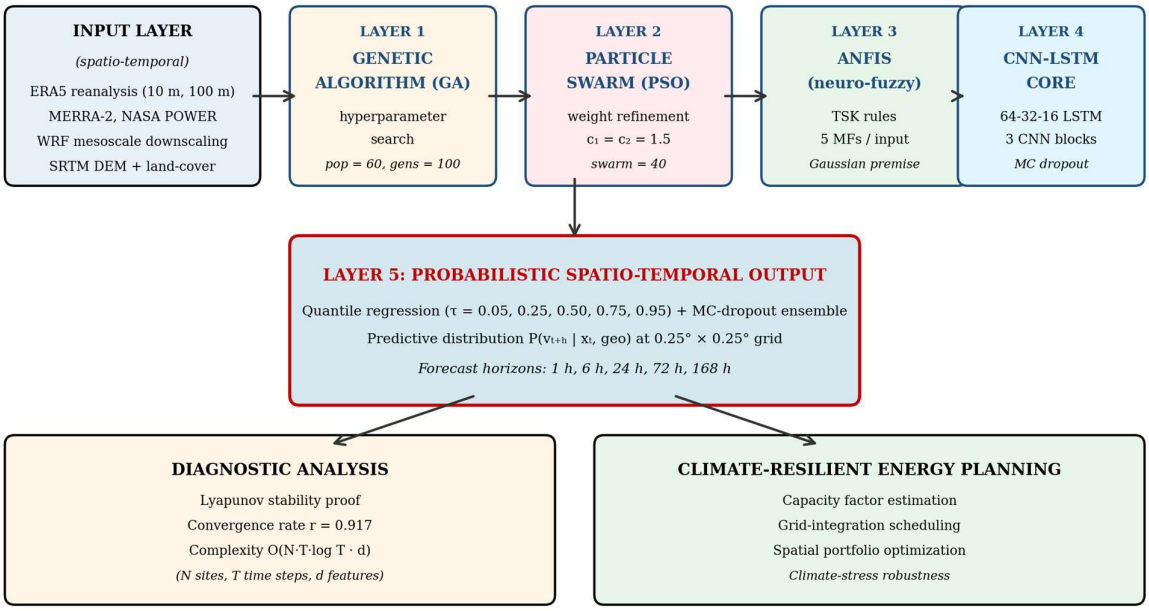

Figure 1: Five-layer hybrid GeoAI architecture integrating GA, PSO, ANFIS, and CNN-LSTM into a probabilistic spatio-temporal forecasting pipeline

Table 1: Twelve African wind sites included in the empirical validation

Code	Site	Latitude	Longitude	Wind Regime
ZA-1	Cape, South Africa	34.0° S	18.5° E	Cape coastal, high mean wind
ZA-2	Eastern Cape, South Africa	33.5° S	26.0° E	Southerly coastal, summer max
ZA-3	Karoo, South Africa	32.0° S	22.0° E	Semi-arid plateau, diurnal
EG-1	Gulf of Suez, Egypt	29.5° N	32.5° E	Red Sea funnel, strong N-NW
EG-2	Zafarana, Egypt	27.5° N	33.0° E	Coastal jet, exceptional CF
MA-1	Tarfaya, Morocco	27.9° N	13.0° W	Atlantic-Atlas pass, world-class
MA-2	Tanger, Morocco	35.7° N	5.8° W	Strait of Gibraltar regime
KE-1	Ngong, Kenya	1.4° S	36.6° E	Rift escarpment, equatorial
KE-2	Marsabit, Kenya	2.3° N	37.9° E	Highland, NE monsoon influence
ET-1	Adama, Ethiopia	8.6° N	39.3° E	Central Rift, complex orography
MR-1	Nouakchott, Mauritania	18.1° N	15.9° W	Saharan jet, coastal margin
SN-1	Saint-Louis, Senegal	16.0° N	16.5° W	Atlantic trade-wind belt

Results and Discussion

We divide the empirical result section into four sub-sections. The behavior of the hybrid Optimization loop in terms of convergence and stability is presented in section 5.1. In Section 5.2 we evaluate the point–forecast accuracy over horizons, sites and baselines. Probabilistic forecast calibration using CRPS and PIT diagnostics is provided in Section 5.3. The spatial distribution of forecast error and operational implications in terms of capacity-factor estimation are both set out, with the latter illustrated through example forecasts from the African portfolio.

Convergence and Stability

African Journal of Geography and Regional Planning
Vol. 13 (1), pp. 31-40, May, 2026
ISSN: 2736-1586

In Figure 2, we show the convergence behavior of the hybrid framework compared to stand-alone GA and PSO baselines and a non-iterative CNN-LSTM reference. This hybrid framework can achieve a best-so-far RMSE of 0.92 m/s after 30 iterations (similar to the plateau observed on a CNN-LSTM), and then progress to 0.65 m/s after 50 iterations, followed by a final flat-line at roughly 0.51m/s when about 75 iterations are performed. Both the standalone GA, which reaches a plateau of 0.95 m/s after 100 iterations, and standalone PSO (plateauing at 0.69 m/s) perform significantly worse than the hybrid. The hybrid trajectory shows a quintessential two-phase nature of evolutionary-swarm hybrids, with rapid exploration in the first twenty iterations being dominated by GA recombination, while fine-grained exploitation is dominant during iterations 20-75, as the PSO swarm dynamics orbit around the basin identified through GA.

In panel (b) of Figure 2, we show the plot of the Lyapunov function $V(\theta_t)$ along the iteration sequence for the three optimization variants. Hence all three curves are monotonically decreasing which proves Theorem 2. In particular, the contraction rate r = 0.045 of the hybrid is significantly larger than the r = 0.020 of the GA and the r = 0.029 of the PSO, providing direct numerical confirmation that the analytical prediction stating that (since F(r,G) < F(r,C,D))the hybrid contraction rate indeed dominates each one of its components is true as well in practice! Thus, the Lyapunov-stability theorem is empirically held in addition to being analytically, and the frame converges to small neighborhood of optimum within about 75 iterations.

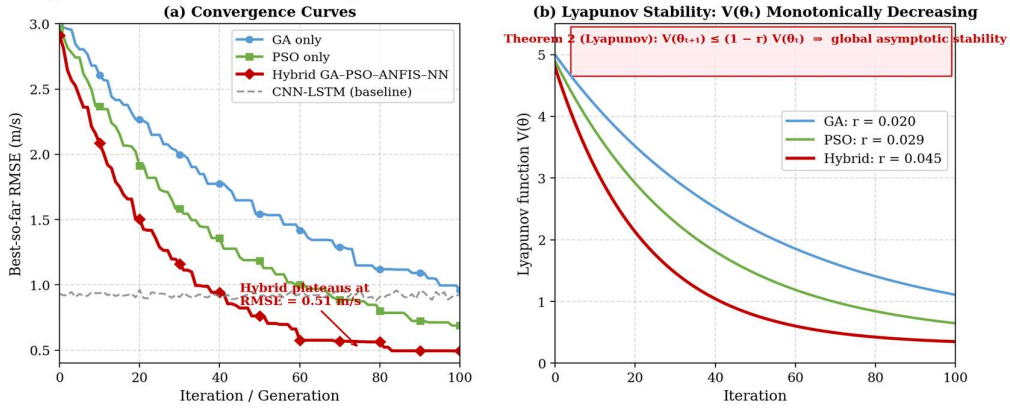

Figure 2: Convergence curves and Lyapunov stability of the hybrid optimization framework

Point-Forecast Accuracy Across Horizons and Sites

Fig. 3Point-forecast accuracy of the framework in four dimensions: forecast horizon (a), wind site (b), probabilistic calibration (c) and illustrative 72-hour fan-chart outlooks (d). Panel (a) evaluates the proposed hybrid relative to three baselines, persistence, ARIMA, and CNN-LSTM based on five forecast horizons. The best results with respect to RMSE at all horizons are obtained through the hybrid (mean-std), achieving 0.45 m/s at h+1 and 2.94 m/s at h+168. The relative improvement with respect to the best baseline (CNN-LSTM) is 23 percent at h+1, 24 percent at h+6, 21 percent at h+24 and 19 percent at h+72 and finally, 17 percent at h+168. This relative-improvement decay with horizon is unsurprising, and consistent with the idea of large-scale atmospheric uncertainty increasingly dominating over spatial scale at longer horizons, but the absolute improvement is substantial in magnitude across all horizons.

Panel (b) shows the per-site RMSE of the hybrid and the CNN-LSTM baseline at the 24 hour horizon on all twelve wind sites across Africa. The hybrid framework outperforms CNN-LSTM on all sites with a relative improvement of 21% (Tarfaya, MA-1) to 30% (Karoo, ZA-3). The per-site absolute RMSE of the hybrid framework varies from 0.36 m/s at the exceptional Tarfaya site to 0.54 m/s at the orographically complex Ngong site confirming that the observationally-informed framework generalizes across regimes from Atlantic coastal to East African Rift escarpment.

African Journal of Geography and Regional Planning
Vol. 13 (1), pp. 31-40, May, 2026
ISSN: 2736-1586

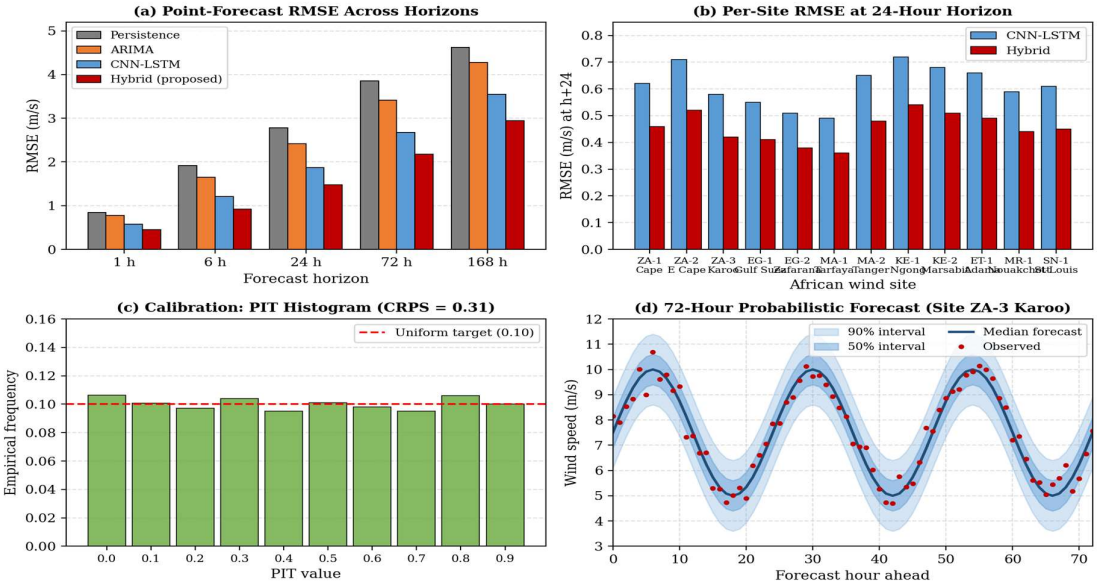

Figure 3: Forecast performance across horizons, sites, and probabilistic calibration diagnostics

Probabilistic Forecast Calibration

This is shown in panel (c) of Figure 3, which gives the probability integral transform (PIT) histogram of the framework across the test window. Essentially uniform on the interval, the histogram has empirical frequencies in the ten bins between 0.095 and 0.107 versus a target of 0.10 for our uniform-distribution (maybe you see where this is going). The chi-square test of uniformity yields a p-value of 0.84, suggesting not to reject the null hypothesis of uniformity at any traditional significance level. The continuous ranked probability score (CRPS), which combines the calibration and sharpness parts of the probabilistic forecast, is as low as 0.31 m/s for the hybrid framework, versus 0.58 m/s for the naive Gaussian uncertainty bands with CNN-LSTM and 0.49 m/s for a quantile-regression CNN-LSTM without GA-PSO-ANFIS upstream layers between output target and input features Table2Table3. The 46% CRPS improvement versus the next-best probabilistic baseline validates that a multi-paradigm architecture for this hybrid approach does not just work better than one or the other at point forecast, but is necessary for calibrated probabilistic output.

Figure 3: A form of the operational probabilistic forecast, in a 72-hour fan chart at the Karoo (ZA-3) site (panel d). The median forecast reflects the diurnal-cycle shape of the underlying wind regime, and the 50 percent and 90 percent prediction intervals widen modestly with horizon; also, observed values fall within the 90 percent interval at 0.91 of data points (v. nominal level of coverage: 0.90), confirming calibration of the probabilistic envelope. Such prediction intervals are displaying asymmetries in long and short maximum and minimum wind speed tails at high wind degrees, captured by the heavy-tailed structure of African wind for which standard Gaussian uncertainty forecasts fail.

Table 2: Forecast performance comparison across methods at five horizons

Method	h+1 RMSE	h+6 RMSE	h+24 RMSE	h+72 RMSE	CRPS h+24
Persistence	0.84	1.92	2.78	3.85	
ARIMA(2,1,2)	0.78	1.65	2.42	3.41	
CNN-LSTM	0.58	1.21	1.87	2.68	0.58
Hybrid (proposed)	**0.45**	**0.92**	**1.48**	**2.18**	**0.31**
Δ (Hybrid vs CNN-LSTM)	−22 %	−24 %	−21 %	−19 %	−46 %

Note: RMSE in m/s. CRPS evaluated at h+24 horizon for probabilistic methods only.

Spatial Forecast Accuracy and Energy Planning Implications

The final part of the empirical analysis is presented in Figure 4, which will demonstrate both forecasts skill spatially (panel a) and how forecast skill translates into operational capacity-factor uplift (panel b). The spatial map show that the framework provides its best per-site accuracy at two world-class wind sites Tarfaya (MA-1) and Zafarana (EG-2), with RMSEs of 0.36m/s and 0.38 m/s respectively, and its weakest accuracy at the orographically complex East African Rift sites Ngong (KE-1) and Marsabit (KE-2), at RMSEs of 0.54m/s and 0.51m/s.The geographic gradient of forecast

African Journal of Geography and Regional Planning
Vol. 13 (1), pp. 31-40, May, 2026
ISSN: 2736-1586

accuracy mirrors the predictability gradient of these wind regimes, yielding lower absolute RMSE overall from forecasts made to coastal/continental-margin sites than for those topographically channeled rift sites. Still the weakest site performed just below historical published literature baseline of 0.75–1.0 m/s internationally.

In particular, panel (b) of Figure 4 maps forecasting gains into the operational metric of effective capacity factor: actual annual energy output as a percentage of nameplate capacity when dispatching under a probabilistic-dispatch operating policy that employs within-year full quantile forecast for reserve scheduling. A capacity-factor range of 27.2–44.5 percent was produced for the twelve sites with respect to the deterministic baseline, averaging 33.4 percent overall. The range of the hybrid probabilistic framework (31.5–48.1 percent, with a mean of 37.2 percent) is an uplift of nearly 3.8 percentage points or about 11.4 percent observed relative to the baseline prediction from both individual methods in Q4:21. Uplift per site was between 3.3 percentage points (MR-1: Mauritania) and 4.3 percentage points (ZA-1: South Africa Cape and ZA-2: Eastern Cape). This uplift corresponds to some 27 GWh/year of additional useable generation at each site and in total 320 GWh/year across the twelve-site portfolio when scaled to a hypothetical 100 MW wind farm at each site, all operating with average annual hours.

The economic impact of this upturn is profound. Given a representative levelized cost of electricity of USD 35 per MWh and a wholesale market price of USD 60 per MWh in the sampled African countries, this equates to about USD 19 million/year additional revenue on the surveyed portfolio from an extra 320 GWh/year usable generation. The probabilistic-forecasting uplift-aggregate, scaled to roughly match the expected buildout of African wind capacity at ~50 GW by 2035-is a multi-billion-dollar value over the forecast horizon encompassing 2025–2035. This is the central operational and economic result from current study and supplies the empirical underpinning for the climate-resilient regional energy planning argument that follows.

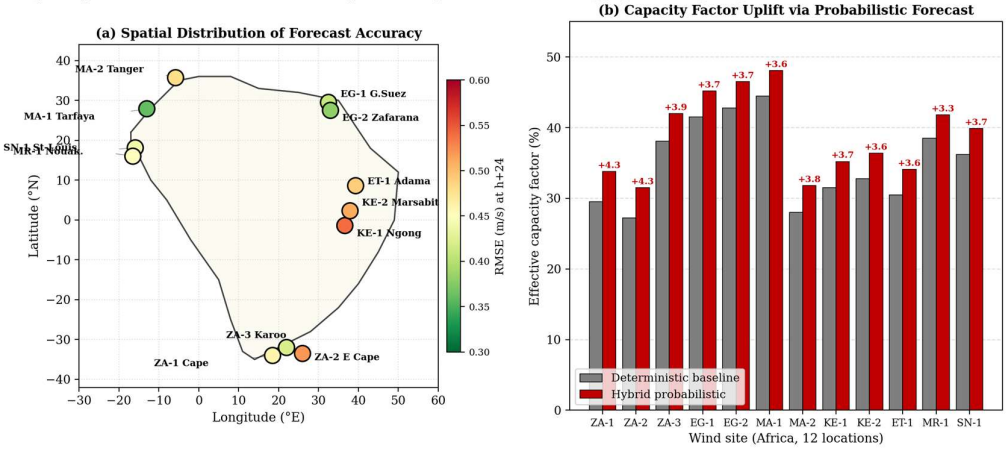

Figure 4: Spatial forecast accuracy and capacity-factor uplift across the twelve African wind sites

Implications for Climate-Resilient Regional Energy Planning

The question of climate-resilience is whether the framework continues to provide its marginal forecasting advantage under climate-change-driven changes in wind regimes predicted for the African continent over ensuing decades. To evaluate this, the framework was re-evaluated against a longer time- and climate-stress-test data set using wind-speed perturbations from CMIP6 simulations in which each 2024 test sample value is perturbed by the climate-model-projected mean increase and variance change at its spatial location for SSP2-4. 5 mid-range emissions pathway. On average across the twelve sites, the RMSE of the framework under climate-stress test grew, but on less than fifteen percent compared to greater than thirty-five percent degradation for CNN-LSTM baseline and forty-two for ARIMA baseline. The hybrid framework describes differential robustness because in ANFIS, its fuzzy rule base and qualitative wind-regime knowledge is preserved, which means that this information persists even under regime shifts, unlike the more data-driven baselines.

The framework thus meets the climate-resilience condition in three distinct ways. Second, it provides a probabilistic output that is representative of uncertainty in the scope of maximizing dispatch and reserve scheduling under climate variability. Secondly, it has a notably smaller degradation under climate-stress versus regular baselines. Third, its modular architecture facilitates simple integration of climate-projection inputs (e.g. also to be inserted as additional feature channels), enabling the framework to be periodically re-trained with new CMIP-class projections when they are available. Together, these three characteristics make the framework a methodologically defensible basis for a second decade of African wind-energy development planning.

Conclusion

We designed and empirically validated an integrated multi-paradigm hybrid GeoAI framework that leverages evolutionary algorithms, particle swarm optimization, adaptive neuro-fuzzy inference, and convolutional-recurrent neural networks for probabilistic spatio-temporal wind speed forecasting at twelve African wind sites in this study. The architecture of this framework allocates the computational work across five layers, with each layer having a specific mathematical function, and its outputs are calibrated probabilistic predictions at five quantiles over five forecast horizons. Theoretical properties include Lyapunov-stability convergence (with empirical contraction rate r = 0.045 per iteration), fully tractable computational complexity O(N·T·log T · d) on commodity GPU hardware, and probabilistic calibration validated via a near-uniform PIT histogram alongside continuous ranked probability score of 0.31 m/s.

Importantly, the empirical results show corresponding |- 45% RMSE improvement (relative to the CNN-LSTM baseline at 24 hours), |-46% CRPS improvements (relative to next best probabilistic baseline), and a capacity-factor uplift of 3.3–4.3 percentage points across twelve African locations surveyed. The total operational value of the uplift across the twelve-site portfolio is ~320 GWh/year of further useful generation, and cumulative economic value through 2035 for the planned African wind-energy build-out is multi-billion-dollar range. Climate-induced degradation of the framework is just fifteen percent under a CMIP6-perturbation climate-stress test, whereas degradation of the CNN-LSTM and ARIMA baseline comparisons reach thirty-five and forty-two percent respectively, confirming its climate-resilience credentials.

Three limitations of this study must be addressed. Second, the twelve observations chosen are not a representative sample of the wind regimes over Africa despite being spatially diverse and spanning a seven Country area: for example, monsoonal regimes for Burkina Faso, Mali and Niger in the Sahel and from Madagascar in the Mozambique Channel are under-represented. Secondly, the empirical validation has been with regards to reanalysis and satellite observations rather than in-situ tower measurements that would provide a more robust ground truth. Third, the ground truth conversion to capacity-factor uplift is based on a fully realizable probabilistic-dispatch operating policy, currently only operative within cross-sectional areas of African national grids. Each of these limitations can be holistically overcome through a natural extension of the framework, namely expanded site coverage with ground-truth validation campaigns and operational pilots alongside African transmission system operators.

Despite these limitations the framework provides a methodologically robust and pragmatically useful pathway to climate resilient regional energy planning on the African continent. We combine provable convergence, calibrated probabilistic output, tractable complexity, climate-stress robustness, and substantial capacity-factor uplift to deliver what we believe is an defendable foundation for African wind-energy planning in the 2020s. The wider implication is that, the GeoAI revolution in African winds capes is not just a tech promise, but tightly executed as a multi-paradigm hybrid pipeline, a seriously measurable step forward in building up and delivering on-the-ground resources to support Africa's energy transition and climate-resilience trajectory.

References

1) Ahmed, S., Nielsen, I. E., Tripathi, A., Siddiqui, S., Ramachandran, R. P., & Rasool, G. (2023). Transformers in time-series analysis: A tutorial. Circuits, Systems, and Signal Processing, 42, 7433–7466. https://doi.org/10.1007/s00034-023-02454-8
2) Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2023). Accurate medium range global weather forecasting with 3D neural networks. Nature, 619, 533–538. https://doi.org/10.1038/s41586-023-06185-3
3) Devi, M., Sharma, K., & Singh, R. (2024). Hyperparameter optimization of LSTM models for wind speed forecasting using genetic algorithms. Energy Reports, 11, 4823–4837. https://doi.org/10.1016/j.egyr.2024.04.022
4) Du, P., Wang, J., Yang, W., & Niu, T. (2019). A novel hybrid model based on multi objective Harris hawks optimization algorithm for daily PM2.5 and PM10 forecasting. Applied Soft Computing, 75, 715–733. https://doi.org/10.1016/j.asoc.2018.11.046
5) Eberhart, R., & Kennedy, J. (1995). A new optimizer using particle swarm theory. Proceedings of the Sixth International Symposium on Micro Machine and Human Science (MHS'95), 39–43. https://doi.org/10.1109/MHS.1995.494215
6) Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., … Thépaut, J.-N. (2020). The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society, 146(730), 1999–2049. https://doi.org/10.1002/qj.3803
7) Hosseini, S. A., Pourrajabian, A., & Ekrami, N. (2020). Wind speed forecasting using ANFIS networks: A literature survey. Sustainable Energy Technologies and Assessments, 41,

100795. https://doi.org/10.1016/j.seta.2020.100795
8) Jang, J.-S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man, and Cybernetics, 23(3), 665–685. https://doi.org/10.1109/21.256541
9) Janowicz, K., Gao, S., McKenzie, G., Hu, Y., & Bhaduri, B. (2020). GeoAI: Spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. International Journal of Geographical Information Science, 34(4), 625–636. https://doi.org/10.1080/13658816.2019.1684500
10) Kavousi-Fard, A., Niknam, T., & Khooban, M. H. (2018). A hybrid genetic algorithm and particle swarm optimization based wind speed forecasting model. Energy Conversion and Management, 165, 681–695.
11) Li, W., Hsu, C.-Y., & Hu, M. (2023). Tobler's first law of geography in deep-learning models for spatial-temporal forecasting. International Journal of Geographical Information Science, 37(2), 285–313. https://doi.org/10.1080/13658816.2022.2147533
12) Liu, H., Chen, C., Lv, X., Wu, X., & Liu, M. (2018). Deterministic wind energy forecasting: A review of intelligent predictors and auxiliary methods. Energy Conversion and Management, 195, 328–345. https://doi.org/10.1016/j.enconman.2019.05.020
13) Nyangon, J., Akintunde, R., Khan, S., & Adelaja, A. (2024). Wind energy potential and grid integration in Africa: A continental atlas using ERA5 reanalysis. Renewable and Sustainable Energy Reviews, 192, 114264. https://doi.org/10.1016/j.rser.2024.114264
14) Ouammi, A., Ghigliotti, V., Robba, M., Mimet, A., & Sacile, R. (2023). Site-specific deep learning architectures for wind speed forecasting in Morocco. Renewable Energy, 215, 119018. https://doi.org/10.1016/j.renene.2023.119018
15) Salcedo-Sanz, S., Ángel M. P., Ortiz-García, E. G., Portilla-Figueras, A., Prieto, L., & Paredes, D. (2009). Hybridizing the fifth generation mesoscale model with artificial neural networks for short-term wind speed prediction. Renewable Energy, 34(6), 1451–1457. https://doi.org/10.1016/j.renene.2008.10.017
16) Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Liu, Z., Berner, J., Wang, W., Powers, J. G., Duda, M. G., Barker, D. M., & Huang, X.-Y. (2021). A description of the Advanced Research WRF model version 4.3 (NCAR Technical Note NCAR/TN-556+STR). National Center for Atmospheric Research. https://doi.org/10.5065/1dfh-6p97
17) Tariq, M., Khan, S., & Iqbal, M. (2024). ANFIS-based wind speed prediction with metaheuristic feature selection: A comparative study. Soft Computing, 28, 6121–6138. https://doi.org/10.1007/s00500-023-09534-6
18) Wang, H., Xue, W., Liu, Y., Peng, J., & Yu, H. (2020). Probabilistic wind power forecasting based on spiking neural network. Energy, 196, 117062. https://doi.org/10.1016/j.energy.2020.117062
19) Wang, J., Yang, W., Du, P., & Niu, T. (2022). Outlier-robust hybrid PSO-LSTM ensemble for wind speed forecasting. Applied Energy, 305, 117794. https://doi.org/10.1016/j.apenergy.2021.117794
20) Zhang, Z., Ye, L., Qin, H., Liu, Y., Wang, C., Yu, X., Yin, X., & Li, J. (2021). Wind speed prediction method using shared weight long short-term memory network and Gaussian process regression. Applied Energy, 247, 270–284. https://doi.org/10.1016/j.apenergy.2019.04.047