Q-SARIMA: A Hybrid Quantum–Classical Extension of SARIMA for Time Series Forecasting in Precision Agriculture

Lucas Grogenski Meloca
Department of Informatics
State University of Maringá (UEM)
Maringá, Brazil
lucas.grogenskimeloca@gmail.com
0000-0001-7057-9327

Rodrigo Clemente Thom de Souza
Advanced Campus in Jandaia do Sul
Federal University of Paraná (UFPR)
Department of Informatics
State University of Maringá (UEM)
Jandaia do Sul, Brazil
thom@ufpr.br
0000-0003-2435-8528

Linnyer Beatrys Ruiz Aylon
Department of Informatics
State University of Maringá (UEM)
Maringá, Brazil
lbruiz@uem.br
0000-0002-4456-6829

Viviana Cocco Mariani
Department of Electrical Engineering
Federal University of Paraná
Curitiba, Brazil
viviana.mariani@ufpr.br
0000-0003-2490-4568

***Abstract*—This paper introduces Q-SARIMA, a hybrid quantum-classical architecture for 10-day meteorological forecasting in precision agriculture. By integrating SARIMA structural robustness with Variational Quantum Circuits (VQC), the model maps qubit expectation values to statistical coefficients using the gradient-free COBYLA optimizer. Evaluation on real-world Brazilian datasets via walk-forward validation shows that while classical SARIMA provides superior stability, Q-SARIMA is competitive for high-volatility series like wind speed. These results highlight NISQ-era challenges, such as convergence instability and cost surface roughness. This study establishes a methodological baseline for hybrid quantum-classical integration, offering critical insights into the trade-offs between quantum expressivity and statistical reliability in agricultural monitoring systems.**

***Index Terms*—Variational Quantum Circuits, Time Series Forecasting, SARIMA, Meteorological Variables, Precision Agriculture**

I. Introduction

Modern precision agriculture relies on accurate meteorological forecasting to optimize resources and mitigate climate risks [1], [2]. While classical statistical methods like Seasonal Autoregressive Integrated Moving Average (SARIMA) offer robustness in modeling seasonal components [3], [4], increasing climate variability often exceeds the generalization capacity of classical Maximum Likelihood Estimators (MLE), necessitating new computational paradigms [5].

This work was supported in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001, in part by the Manna Team, and by Softex, Fundação Araucária, and Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) under grant 402015/2025-8 (chamada CNPq/MCTI No. 44/2024 - Faixa B).

In this context, Quantum Machine Learning (QML) via Variational Quantum Algorithms (VQA) has emerged as a disruptive approach for exploring complex optimization landscapes in the Noisy Intermediate-Scale Quantum (NISQ) era [6]. However, a significant gap remains in the direct hybridization of classical stochastic models with Variational Quantum Circuit (VQC). Most current research replaces classical components entirely, often neglecting the proven structural efficacy of traditional time series decomposition.

This paper proposes Quantum Seasonal Autoregressive Integrated Moving Average (Q-SARIMA), a hybrid architecture that integrates classical SARIMA structures with variational quantum optimization. The core hypothesis is that a VQC, acting as a parameter optimizer via Constrained Optimization by Linear Approximations (COBYLA), can explore the solution space more effectively than deterministic methods, potentially capturing subtle stochastic correlations in variables like wind speed [5], [7].

The methodology establishes a computational pipeline using PennyLane [8] to map qubit parameters to model coefficients. We evaluate Q-SARIMA against classical SARIMA using real-world 10-day meteorological data from distinct Brazilian climatic regimes (Am and Aw) [9]. Model performance is quantified through Root Mean Square Error (RMSE), Adjusted Coefficient of Determination (R^2 adj.), and Mean Bias Error (MBE).

In alignment with open science principles, the implementation code and datasets utilized in this study are openly available on Zenodo at: https://doi.org/10.5281/zenodo.19446212.

II. Methodology

This section describes the hybrid Q-SARIMA framework for agro-meteorological forecasting, integrating classical structural identification with VQAs for parameter optimization in the Hilbert space while maintaining SARIMA interpretability.

A. Data Acquisition and Preprocessing

The performance of data-driven models in precision agriculture strongly depends on the quality of input features. This study employs meteorological data obtained from the National Aeronautics and Space Administration/Prediction of World Wide Energy Resources (NASA/POWER) project [5], [10], accessed via Application Programming Interfaces (API). The dataset spans from January 1, 2003, to December 31, 2019.

1) Variable Selection and Aggregation: Following the methodology of [9], four variables were selected: air temperature at 2 meters (°C), relative humidity (%), wind speed at 2 meters ($\mathrm{m/s}$), and all-sky surface shortwave downward irradiance ($\mathrm{MJ/m^2/day}$). Data were collected for fifteen major Brazilian grain-producing municipalities and grouped according to their predominant Köppen climate classification [11]. Data were collected for fifteen major Brazilian grain-producing municipalities, grouped into two representative time series according to their predominant Köppen climate classification: Tropical Monsoon (Am, 5 municipalities) and Tropical Savanna (Aw, 10 municipalities). Daily means were computed across locations to capture regional climatological behavior.

2) Temporal Structuring: In accordance with agrometeorological practices for crop monitoring [9], daily observations were aggregated into 10-day periods (deciles). This temporal resolution aligns with crop phenological stages while reducing high-frequency noise irrelevant to strategic decision-making. As a result, the transformed series exhibit a seasonal periodicity of $m = 36$ observations per year. The final dataset consists of eight univariate series, corresponding to four variables under two climatic regimes, enabling robustness analysis.

3) Stationarity and Partitioning: The SARIMA framework requires stationarity; therefore, the non-seasonal (d) and seasonal (D) differencing orders are automatically identified during the model selection stage using statistical unit root tests, such as the Augmented Dickey-Fuller (ADF) test.

Temporal data partitioning is adopted to avoid information leakage. Instead of random shuffling, a walk-forward validation scheme is applied. For each test year, the model is trained using historical data starting in 2003. A hybrid windowing strategy is employed: the training window expands until reaching a three-year span and subsequently advances as a sliding window with fixed length, reflecting a realistic operational forecasting scenario.

B. The Hybrid Q-SARIMA Architecture

The Q-SARIMA model hybridizes the Box-Jenkins methodology with QML through a sequential two-phase pipeline (Fig. 1).

1) Phase 1: Classical Structure Discovery: The SARIMA order $(p, d, q) \times (P, D, Q)_m$ was dynamically identified using the `auto_arima` algorithm [12] by minimizing the Akaike Information Criterion (AIC).

The selected structure defines the specific polynomial form of the model:

$$\Phi_P(B^m)\phi_p(B)(1-B)^d(1-B^m)^D y_t = \Theta_Q(B^m)\theta_q(B)\varepsilon_t \,, \tag{1}$$

where B is the lag operator, y_t is the observed value at time t, ε_t is the white noise error term, ϕ_p and θ_q are the non-seasonal autoregressive and moving average polynomials of orders p and q, respectively, and Φ_P and Θ_Q are their seasonal counterparts of orders P and Q. The total number of coefficients to be estimated is given by $k = p+q+P+Q$. To ensure the feasibility of quantum simulation on NISQ devices and prevent overfitting on limited data, a constraint is imposed such that $k \leq 4$. This effectively limits the size of the quantum ansatz to a maximum of 4 qubits.

2) Phase 2: Variational Quantum Optimization: Parameter estimation is delegated to a VQC. We employ a Hardware-Efficient Ansatz [13], detailed within the optimization loop of Fig. 1, in PennyLane [8], where the qubit count equals the number of coefficients k identified in Phase 1.

The coefficients $\boldsymbol{\beta}$ are extracted via Pauli-Z expectation values: $\beta_i = \langle\psi(\boldsymbol{\theta})|\hat{Z}_i|\psi(\boldsymbol{\theta})\rangle$. Here, $|\psi(\boldsymbol{\theta})\rangle$ is the parameterized quantum state generated by the VQC, $\boldsymbol{\theta}$ represents the vector of trainable rotation angles (variational parameters), and $\hat{Z}_i$ is the Pauli-Z observable applied to the i-th qubit. While mapping coefficients to the $[-1, 1]$ interval intuitively supports stability for first-order AR or MA processes, it does not strictly guarantee that the roots of the characteristic polynomial lie outside the unit circle for higher-order models. However, in practice, this bounded search space functions as a heuristic regularization mechanism. It prevents the optimizer from exploring highly divergent parameter regions, effectively penalizing explosive processes and encouraging model convergence within a localized, stable regime. The optimization loop minimizes the in-sample Mean Squared Error (MSE) calculated in the original data scale:

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{T'}\sum_{t=1}^{T'}(y'_t - \hat{y}'_t(\boldsymbol{\beta}(\boldsymbol{\theta})))^2 \,, \tag{2}$$

where T' is the length of the in-sample training window, and y'_t and $\hat{y}'_t$ denote the true and predicted values of the stationary (differenced) time series at step t, respectively. The gradient-free COBYLA optimizer [14] manages the feedback loop over a maximum of 50 iterations, a strategy chosen for its robustness against NISQ-era simulation noise. To ensure evaluation rigor, parameters $\boldsymbol{\theta}$ undergo a random "cold start" initialization at each step of the walk-forward validation protocol.

C. Forecasting and Reconstruction

Once the optimal parameters $\boldsymbol{\theta}^*$ are converged upon, the final coefficients $\boldsymbol{\beta}^*$ are extracted. The model then generates

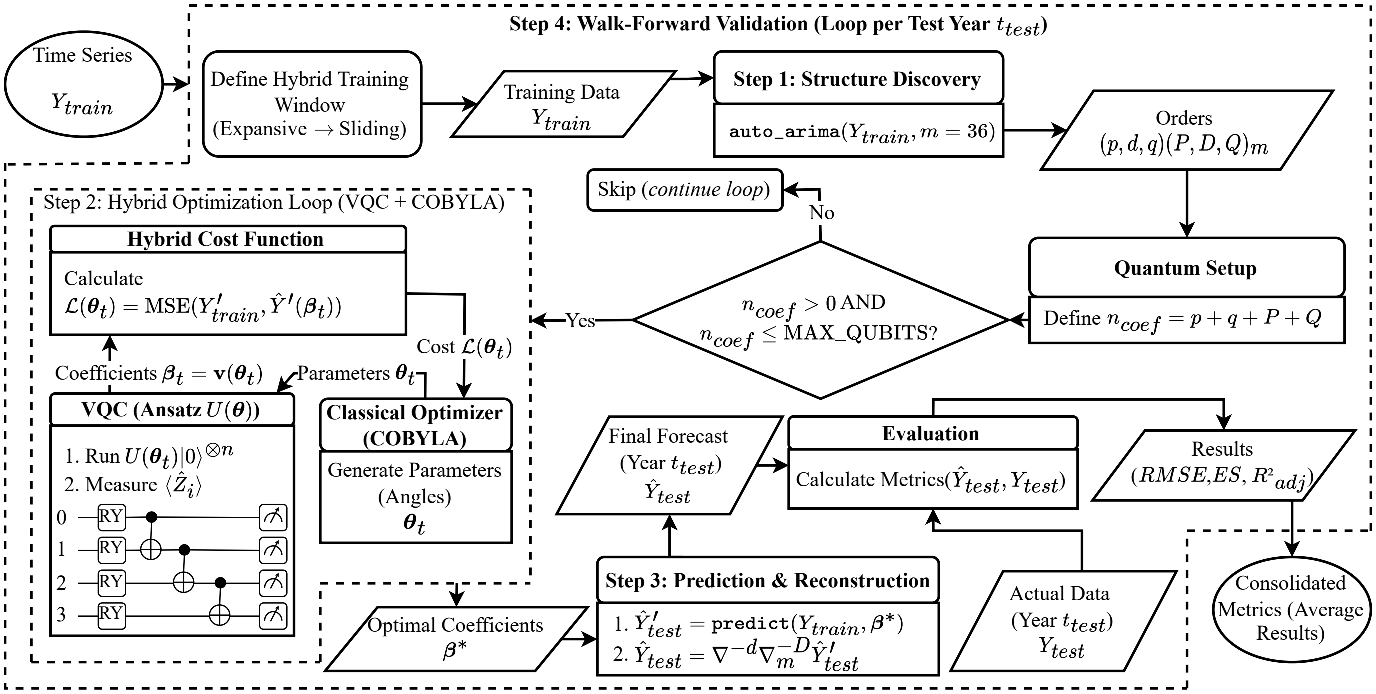

Fig. 1. Overall architecture of the Q-SARIMA methodological pipeline. The schematic illustrates the hybrid quantum-classical optimization loop managed by the COBYLA optimizer, detailing the internal Hardware-Efficient VQC architecture composed of parameterized rotation gates (R_y) and linear entanglement layers (CNOT) for a 4-qubit configuration.

an out-of-sample forecast for the horizon $h = 1$ (one decile). Since the internal model operates on differenced data to satisfy stationarity, a reconstruction step is necessary. This involves applying the inverse difference operators ∇^{-d} and ∇_m^{-D} (cumulative summation) to the forecasted values, restoring the prediction to the original scale of the meteorological variable.

D. Experimental Design and Evaluation

1) Walk-Forward Validation: To rigorously assess performance, we implement a hybrid expanding-then-sliding walk-forward validation protocol. For every test year, the model structure and parameters are dynamically updated using available historical data. This approach ensures that the forecasting system adapts to non-stationary climate shifts and provides a realistic estimation of model accuracy in operational agricultural environments.

2) Evaluation Metrics: The predictive performance is quantified using three metrics, aligned with standard practices in agricultural data analytics [9], [15], [16]:

- **RMSE:** Measures the average magnitude of the error in the units of the variable:

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N}(y_t - \hat{y}_t)^2}. \tag{3}$$

- $\mathbf{R}^2_{\text{adj}}$: Evaluates the proportion of variance explained by the model, penalized by the number of estimated parameters (k). This allows for a fair comparison between models of varying complexity:

$$R^2_{\text{adj}} = 1 - \left(\frac{(1 - R^2)(N - 1)}{N - k - 1}\right), \tag{4}$$

where N is the total number of out-of-sample observations and R^2 is the standard coefficient of determination.

- **MBE:** Also referred to as Systematic Error (SE), it quantifies the average tendency of the model to over- or underestimate the observed values:

$$\text{MBE} = \frac{1}{N}\sum_{t=1}^{N}(\hat{y}_t - y_t). \tag{5}$$

3) Computational Viability: Given the computational overhead of simulating quantum circuits, we also evaluate the execution time. Experiments were conducted on a workstation equipped with an AMD Ryzen 5 5600G processor and an AMD Radeon RX 6600 GPU. The analysis focuses on the trade-off between the potential accuracy gains offered by the quantum exploration of the parameter space and the increased computational cost relative to classical MLE methods.

III. Results and Discussion

This section compares the hybrid Q-SARIMA (QSAR) and classical SARIMA (SAR) models across Am and Aw climate regimes. Detailed RMSE, R^2 adj., and MBE metrics are provided in Table I and Table II.

For mean temperature (T2M), classical SARIMA generally provided superior stability. In the Am regime, SARIMA's

TABLE I

COMPARISON OF ERROR METRICS (RMSE, R^2 ADJ., MBE) SPECIFICALLY FOR THE AM CLIMATE REGIME ACROSS ALL METEOROLOGICAL VARIABLES. THE ALLSKY_SFC_SW_DWN VARIABLE WAS ABBREVIATED AS ALLSKY.

			T2M		RH2M		WS2M		ALLSKY	
In	**Out**	**Metric**	**QSAR**	**SAR**	**QSAR**	**SAR**	**QSAR**	**SAR**	**QSAR**	**SAR**
5	13	RMSE	2.410	1.842	6.276	5.050	0.581	0.213	7.340	2.631
		Adj. R²	-0.160	0.327	-0.300	0.108	-6.369	-0.064	-6.635	-0.053
		MBE	0.546	0.636	0.066	-1.729	-0.334	0.063	1.581	1.242
8	10	RMSE	2.586	2.063	6.304	4.556	0.311	0.199	23.757	2.750
		Adj. R²	-0.381	0.030	-0.277	0.277	-1.395	0.045	-129.2	-0.527
		MBE	0.609	1.073	0.683	-0.732	-0.184	0.011	-2.999	1.641
10	8	RMSE	8.316	2.285	6.156	5.160	0.361	0.202	13.316	2.880
		Adj. R²	-17.51	-0.422	-0.177	0.111	-2.235	-0.101	-46.08	-1.023
		MBE	-1.752	1.290	1.711	-0.259	-0.263	-0.031	-4.911	1.916
12	6	RMSE	7.536	2.411	6.967	5.507	0.478	0.202	5.374	2.883
		Adj. R²	-21.32	-1.405	-0.758	0.013	-6.013	-0.351	-9.842	-1.921
		MBE	-3.113	1.505	1.100	-0.723	-0.208	-0.046	0.352	2.158
14	4	RMSE	2.318	2.070	7.852	4.984	0.383	0.201	6.401	2.580
		Adj. R²	-2.177	-1.450	-0.946	0.227	-4.181	-0.542	-16.77	-2.051
		MBE	1.055	1.131	2.108	0.333	-0.168	-0.072	-3.595	1.941
16	2	RMSE	8.405	1.488	7.715	5.900	0.183	0.190	5.210	1.792
		Adj. R²	-67.28	-0.703	-1.066	-0.092	-0.196	-0.413	-28.91	-0.885
		MBE	-4.127	0.300	3.198	3.215	-0.085	-0.039	-3.462	0.660

TABLE II

COMPARISON OF ERROR METRICS (RMSE, R^2 ADJ., MBE) SPECIFICALLY FOR THE AW CLIMATE REGIME ACROSS ALL METEOROLOGICAL VARIABLES. THE ALLSKY_SFC_SW_DWN VARIABLE WAS ABBREVIATED AS ALLSKY.

			T2M		RH2M		WS2M		ALLSKY	
In	**Out**	**Metric**	**QSAR**	**SAR**	**QSAR**	**SAR**	**QSAR**	**SAR**	**QSAR**	**SAR**
5	13	RMSE	1.138	1.181	80.516	5.822	0.514	0.134	5.780	1.847
		Adj. R²	0.107	0.143	-150.6	0.470	-12.73	0.042	-10.89	-0.262
		MBE	0.164	0.423	-12.29	-1.153	0.026	0.044	-2.277	0.741
8	10	RMSE	7.571	1.315	10.011	6.762	0.287	0.119	5.247	1.842
		Adj. R²	-39.10	-0.025	-0.365	0.344	-2.542	0.283	-10.70	-0.462
		MBE	-2.173	0.566	-2.541	1.211	-0.178	0.007	-2.638	0.974
10	8	RMSE	5.585	1.385	8.302	6.798	0.264	0.118	7.239	1.832
		Adj. R²	-23.81	-0.189	0.086	0.302	-2.682	0.137	-20.62	-0.505
		MBE	-1.625	0.807	2.687	3.127	-0.212	-0.031	-3.558	1.014
12	6	RMSE	4.044	1.499	24.304	7.079	0.316	0.121	7.749	1.877
		Adj. R²	-11.99	-0.404	-9.522	0.175	-8.607	-0.260	-33.78	-1.150
		MBE	-1.828	0.772	-3.122	3.702	-0.215	-0.047	-4.141	1.339
14	4	RMSE	6.222	1.118	21.161	6.206	0.264	0.127	5.061	1.721
		Adj. R²	-28.95	-0.021	-6.953	0.358	-6.078	-0.565	-25.61	-2.181
		MBE	-2.320	0.482	-4.646	3.021	-0.189	-0.057	-1.531	1.410
16	2	RMSE	1.731	0.985	18.060	6.523	0.172	0.143	5.126	1.270
		Adj. R²	-5.834	-0.251	-5.447	0.163	-1.980	-0.984	-33.77	-1.020
		MBE	-0.272	0.311	-4.295	4.278	-0.060	-0.053	-2.704	0.838

RMSE ranged from 1.488 to 2.285. Conversely, Q-SARIMA exhibited higher errors (up to 8.405) and deeply negative R^2 adj. values (e.g., -67.28 in the 16-2 setup), consistently underestimating temperatures in wider windows (MBE -4.127). SARIMA maintained a persistent positive bias (MBE up to 1.505). However, Q-SARIMA proved competitive in the Aw 5-13 configuration, achieving an RMSE of 1.138 and a lower MBE (0.164) than the classical model.

Relative humidity (RH2M) highlighted convergence difficulties for the hybrid architecture. In Aw 5-13, Q-SARIMA produced an extreme RMSE of 80.516, an R^2 adj. of -150.6, and severe underestimation (MBE -12.29). This catastrophic divergence suggests the COBYLA optimizer struggles with local minima or barren plateaus in high-variance savanna climates. In the Am regime, Q-SARIMA errors were smaller but still lagged behind the baseline (e.g., RMSE 6.156 vs. 5.160 in the 10-8 setup).

Conversely, wind speed (WS2M) yielded the hybrid model's most favorable results, remaining highly competitive across both climates. In Am 16-2, Q-SARIMA's RMSE of 0.183 marginally outperformed SARIMA's 0.190. As Fig. 2(a) illustrates, Q-SARIMA successfully tracked high-volatility peaks in 2015 that the classical model missed. Fig. 2(b) confirms this consistency, with quantum predictions clustering tightly along the identity line. Even in the Aw regime (16-2), Q-SARIMA's RMSE (0.172) closely matched SARIMA (0.143).

Global solar radiation (ALLSKY) analysis confirms a performance gap in the current hybrid implementation's ability to model strong seasonality. SARIMA remained robust with RMSE values between 1.270 and 2.883 and a tendency toward slight overestimation (MBE up to 2.158). Q-SARIMA, however, showed significant instability, with RMSE spiking to

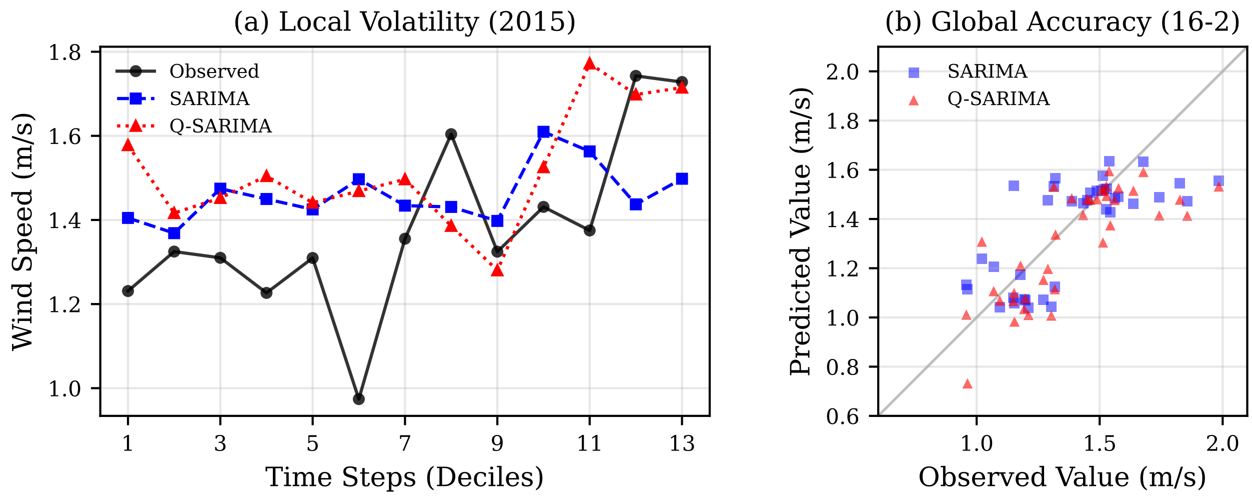

Fig. 2. Visual performance analysis of Q-SARIMA vs. SARIMA for Wind Speed (WS2M). (a) The forecast trajectory for the year 2015 (Am regime, 5-13 configuration) highlights the Q-SARIMA model's ability to track high-frequency volatility peaks. (b) The global scatter plot for the 16-2 configuration demonstrates the hybrid model's consistency across the full dataset, showing tight clustering along the diagonal identity line.

23.757 in the Am 8-10 configuration and frequently failing to reach observed solar peaks (MBE of -4.911 in the Am 10-8 setup). The consistently negative R^2 adj. values suggest that the hardware-efficient ansatz lacks the expressivity required for these highly atmospheric and seasonal series.

In summary, while the hybrid Q-SARIMA model demonstrates technical feasibility and excels in high-volatility scenarios like wind speed, it lacks the overall robustness of classical MLE for general agricultural monitoring. The observed instabilities reflect NISQ-era challenges, including variational cost surface roughness and gradient-free optimization limits in high-dimensional spaces. These results establish a necessary baseline for future refinements in quantum-classical integration for environmental modeling.

A. *Observed Patterns and Limitations*

The analysis reveals a heterogeneous performance landscape for the hybrid Q-SARIMA architecture. A recurring pattern is the higher volatility of the quantum optimizer compared to the deterministic MLE. While the classical SARIMA exhibits stable convergence, Q-SARIMA demonstrates significant performance fluctuations as the input window increases. This behavior links to VQC sensitivity to initialization and the complex cost surface, resulting in unstable MBE values fluctuating between overestimation and underestimation across training windows.

Climate regime also influences robustness. In the stable Am regime, the hybrid model maintains relative stability, whereas the Aw regime (marked by abrupt seasonal transitions) caused convergence failures. Notably, relative humidity forecasting showed massive RMSE spikes and systematic underestimation (MBE -12.29). These results suggest the Hardware-Efficient Ansatz lacks the inductive bias for non-linear seasonal shifts. Furthermore, performance degradation in longer horizons (e.g., 10-8, 12-6) highlights the "barren plateau" phenomenon, where vanishing gradients trap the COBYLA optimizer in suboptimal regions, leading to high bias in variables like mean temperature.

Experimental limitations are primarily dictated by the constraints of the NISQ era. The current 4-qubit architecture, while necessary for simulation feasibility, strictly bounds the SARIMA structural capacity. This restriction hinders the model's ability to fully capture complex seasonal patterns. Furthermore, the reliance on a baseline Hardware-Efficient Ansatz with single R_y rotations and a single-variable mapping limits the expressivity of the variational circuit and does not fully exploit quantum entanglement. Consequently, while the current hybrid model serves as a functional proof-of-concept, its scalability to high-precision operational environments depends on exploring deeper, problem-specific ansatzes as quantum hardware scales.

TABLE III
COMPARISON OF THE AVERAGE COMPUTATIONAL EXECUTION TIME (IN SECONDS) BETWEEN THE SARIMA AND Q-SARIMA MODELS FOR THE DIFFERENT METEOROLOGICAL VARIABLES AND CLIMATE REGIMES ANALYZED. THE ALLSKY_SFC_SW_DWN VARIABLE WAS ABBREVIATED AS ALLSKY IN THIS TABLE.

		QSARIMA	SARIMA
Am	T2M	152.48	118.00
	WS2M	52.95	59.03
	RH2M	88.14	36.97
	ALLSKY	161.49	79.89
Aw	T2M	81.63	132.49
	WS2M	66.28	99.44
	RH2M	71.51	59.67
	ALLSKY	180.72	13.03

B. *Computational Cost Analysis*

Computational overhead is critical for viability. Table III summarizes execution times.

In general, the Q-SARIMA model imposes a significant computational burden, with execution times for the ALLSKY variable in the Aw regime reaching 180.72 seconds, over 13 times the duration of the classical SARIMA (13.03s). This discrepancy is attributed to the iterative nature of the VQC-COBYLA loop, which requires repeated state preparation and measurement to evaluate the cost function. However, a notable exception is observed in the T2M variable for the Aw regime, where the Q-SARIMA converged in 81.63 seconds compared to 132.49 seconds for the classical MLE. This counterintuitive result suggests that gradient-free optimizers like COBYLA may find satisfactory local minima more rapidly than gradient-based classical methods in specific noisy or flat likelihood surfaces, even if the resulting parameters are less precise.

Scaling remains a critical challenge. As the number of qubits increases to match higher-order SARIMA structures, the classical simulation of the quantum state scales exponentially, leading to non-linear increases in training time. For practical deployment in edge-computing agricultural sensors, the current hybrid overhead is prohibitive. Future scalability depends on migrating from classical simulation to actual Quantum Processing Units (QPUs), where the intrinsic parallelism of the hardware could theoretically mitigate the temporal costs of evaluating complex variational circuits.

C. Scientific Discussion

Integration reveals complex trade-offs; while viable, Q-SARIMA lacks SARIMA's consistency. Classical MLE remains superior for linear seasonality (e.g., solar radiation) due to stable convergence and predictable bias profiles.

WS2M competitiveness suggests quantum probability and non-convex search suit volatile, low-autocorrelation series. Q-SARIMA's competitive RMSE (e.g., 0.183 in Am 16-2) indicates it captures nuances where MLE converges prematurely. Unlike black-box QNNs, Q-SARIMA preserves coefficient identity, ensuring structural interpretability via statistical diagnostics.

NISQ-era noise and barren plateaus drive observed instability. The unstructured Hardware-Efficient Ansatz complicates training and induces systematic errors. Successful hybrid forecasting requires balancing quantum expressivity with structural constraints to ensure stability and bias control.

IV. Conclusion

This paper evaluated Q-SARIMA, a hybrid architecture integrating SARIMA robustness with VQC optimization for meteorological forecasting in precision agriculture. Empirical results show that while classical SARIMA remains superior in stability and accuracy, the hybrid model demonstrates localized potential in temperature forecasting. Its volatility in stochastic variables like humidity underscores current NISQ-era challenges, particularly regarding noise sensitivity in non-linear optimization landscapes.

The primary contribution is the validation of a variational framework that maps Hilbert space representations to interpretable statistical parameters. This proof-of-concept establishes that quantum circuits can function as stochastic optimizers within established Box-Jenkins methodologies, providing a foundation for leveraging non-classical correlations in data-driven agricultural systems. Future efforts will focus on general methodological refinements and validation across diverse geographical datasets to enhance the robustness of agricultural forecasting.

References

[1] M. O'Grady, D. Langton, F. Salinari, P. Daly, and G. O'Hare, "Service design for climate-smart agriculture," *Information Processing in Agriculture*, vol. 8, no. 2, pp. 328–340, 2021.

[2] H. Hadjur, D. Ammar, and L. Lefèvre, "Toward an intelligent and efficient beehive: A survey of precision beekeeping systems and services," *Computers and Electronics in Agriculture*, vol. 192, p. 106604, 2022.

[3] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Hoboken, NJ: John Wiley & Sons, 2015, updated edition including modern approaches to ARIMA and state-space modeling.

[4] K. B. Banakara, N. Sharma, S. Sahoo, S. K. Dubey, and V. M. Chowdary, "Evaluation of weather parameter-based pre-harvest yield forecast models for wheat crop: a case study in saurashtra region of gujarat," *Environmental Monitoring and Assessment*, vol. 195, no. 1, p. 51, 2022.

[5] S. Shukla and S. Pasari, "Short-term wind speed prediction with adaptive signal processing based hybrid statistical models," *Energy Systems*, 2025.

[6] S. M. Tripathi, H. Upadhyay, and J. Soni, "Analysis of ansätze expressibility and complexity and their impact on classification accuracy using QNN and QLSTM," *IEEE Access*, vol. 13, pp. 152 412–152 429, 2025.

[7] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.

[8] V. Bergholm *et al.*, "Pennylane: Automatic differentiation of hybrid quantum-classical computations," 2022. [Online]. Available: https://arxiv.org/abs/1811.04968

[9] V. Vieira, J. Silva, and R. Santos, "Spatial and temporal analysis of agricultural productivity," *Revista Brasileira de Geografia Física*, vol. 17, no. 3, pp. 1434–1456, 2024.

[10] F. Zaka, I. A. Nafisah, J. Lin, M. M. A. Almazah, I. Hussain, M. Almazroui, R. Chakrabortty, and H. Louati, "Real-time temperature nowcasting using deep learning models across multiple locations," *Modeling Earth Systems and Environment*, vol. 11, no. 3, p. 165, 2025.

[11] C. A. Alvares, J. L. Stape, P. C. Sentelhas, J. d. M. Gonçalves, G. Sparovek *et al.*, "Köppen's climate classification map for Brazil," *Meteorologische zeitschrift*, vol. 22, no. 6, pp. 711–728, 2013.

[12] T. G. Smith *et al.*, "pmdarima: Arima estimators for Python," 2017. [Online]. Available: http://www.alkaline-ml.com/pmdarima

[13] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.

[14] M. J. D. Powell, *A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation*. Dordrecht: Springer Netherlands, 1994, pp. 51–67.

[15] C. J. Willmott, "Some comments on the evaluation of model performance," *Bulletin of the American Meteorological Society*, vol. 63, no. 11, pp. 1309–1313, 1982.

[16] J. Cornell and R. Berger, "Factors that influence the value of the coefficient of determination in simple linear and nonlinear regression models," *Phytopathology*, vol. 77, no. 1, pp. 63–70, 1987.