Quantum-Enhanced Chaotic Particle Swarm Optimization for LSTM-Based Stock Market Forecasting

Subhas Mondal[1][0009−0003−1946−8906], Saptadeep Biswas[1,2][0000−0002−2092−5487], Seyedali Mirjalili[3,4][0000−0002−1443−9458], Absalom E. Ezugwu[5][0000−0002−3721−3400], and Uttam Kumar Bera[1][0000−0001−5426−7614]

[1] Department of Mathematics, National Institute of Technology Agartala, Tripura, India subhas_24odmar001@phd.nita.ac.in, uttam.math@faculty.nita.ac.in

[2] Department of Operations and Supply Chain, GITAM School of Business, GITAM (Deemed to be University), Visakhapatnam, India sbiswas3@gitam.edu, saptadeepmath.sch@nita.ac.in

[3] Centre for Artificial Intelligence Research and Optimisation, Torrens University Australia, Brisbane, Australia

[4] University Research and Innovation Center, Obuda University, 1034, Budapest, Hungary ali.mirjalili@torrens.edu.au

[5] Unit for Data Science and Computing, North-West University, 11 Hoffman Street, Potchefstroom 2520, South Africa absalom.ezugwu@nwu.ac.za

Abstract. Financial market prediction is challenging due to inherent volatility and nonlinear price dependencies. This paper introduces the Hybrid Quantum Chaotic Particle Swarm Optimization (HQCPSO) algorithm, which integrates quantum-inspired probabilistic position updates derived from δ potential well mechanics with chaotic mapping sequences to optimize Long Short-Term Memory (LSTM) network hyperparameters for stock price forecasting. Unlike classical Particle Swarm Optimization (PSO), HQCPSO eliminates velocity tracking and instead governs particle movement through quantum wave-function probability densities, while chaotic attractors replace fixed random coefficients to sustain search diversity. Validation on the CEC-2022 benchmark suite demonstrates rank-1 performance (mean rank 1.5) against PSO (3.67), Grasshopper Optimization Algorithm (GOA, 2.83), Grey Wolf Optimizer (GWO, 2.67), and Musk Ox Optimizer (MO, 4.33) across 12 test functions, with Wilcoxon rank-sum significance confirmed $p < 0.05$ against PSO, GWO, and MO. Applied to five years of National Stock Exchange (NSE) of India data for four equity indices, HQCPSO-optimized LSTM achieves a 27.7% Mean Absolute Error (MAE) reduction for Reliance and 49.6% Mean Squared Error (MSE) reduction for Nestle India relative to standard PSO, with R^2 scores of 0.916–0.941. The quantum-chaotic framework maintains 3–5% residual exploration in late iterations, effectively balancing exploration and exploitation. These results validate HQCPSO's utility for robust hyperparameter optimization in real-world financial forecasting.

Keywords: Quantum-Chaotic PSO · LSTM Hyperparameter Optimization · Stock Market Prediction · Metaheuristics · Financial Time Series· Swarm Intelligence · Deep Learning

1 Introduction

Stock markets are important economic indicators that reflect a complex interplay of macroeconomic conditions, investor sentiment, political developments, and company fundamentals. This makes accurate price forecasting a high-value problem, with the caveat that financial time series are inherently nonlinear and non-stationary [5].

Deep learning has become a more fruitful approach for this task. Among these architectures, Long Short-Term Memory (LSTM) networks are especially well-suited for temporal sequence modeling: their gating mechanisms dynamically retain or discard information across time steps, which helps to overcome the vanishing-gradient problem that hinders standard recurrent networks and solves both short-term fluctuations and long-range dependencies in price data [7].

However, LSTM model predictive power is extremely sensitive to hyperparameter choices, layer depth, neuron counts, dropout rates, and learning schedules that define a high-dimensional, non-convex search space. Exhaustive search methods, including grid search and random search, become computationally intractable at scale [4]. For example, Particle Swarm Optimization (PSO) [3] is a population-based alternative to gradient-based methods, whereby each candidate solution is part of a swarm that moves toward the personal best and global best positions guided by velocities. While standard PSO is efficient on many benchmark problems, it is known to suffer from premature convergence, especially among multimodal landscapes[10], and when the fitness surface is non-stationary, which also tends to happen in hyperparameter tuning for deep learning[20].

To mitigate premature convergence, quantum-inspired optimization uses probabilistic position sampling from wells of a potential well associated with the quantum mechanics of an imaginary particle to replace deterministic velocity updates and broaden the stochastic search without explicit momentum terms[12]. In a separate study, it was shown that chaotic mapping adaptation improves population diversity by replacing pseudo-random sequences with deterministic but ergodic chaotic iterates, which ensures more dense and uniform coverage across the search space[11]. Although both of these concepts improve on classical PSO in isolation, their effect when combined on high-dimensional deep learning hyperparameter problems has not been studied in a systematic way.

This paper fills this gap through the following contributions:

- HQCPSO: A novel hybrid approach combining δ potential-well quantum position update with a chaotic attractor sequence for improved adaptability in exploring-exploiting trade-offs.
- Thorough CEC-2022 benchmark testing with statistical endorsement (Wilcoxon rank-sum tests) against PSO, GOA[13], GWO [14], and MO technique [19],

including progress towards optima dynamics analysis and exploration-exploitation ratios.
- Application to hyperparameter optimization of LSTM for four NSE India equity indices across various market sectors, with multi-metric performance reporting and computational complexity evaluation

2 Literature Review

2.1 Deep Learning for Financial Forecasting

Hybrid deep learning approaches were demonstrably successful in recent work in computational finance. Chen et al. [6] applied sentiment analysis to an anticipatory model for predicting trends, and Gulmez [8] introduced GA-Attention-Fuzzy-Stock-Net that integrated genetic algorithms with attention networks and fuzzy logic. Rezaei et al. [16], have applied complementary ensemble empirical mode decomposition (CEEMD) and empirical mode decomposition (EMD) for preprocessing of the input signal prior to feature extraction using CNN-LSTM, achieving improvements upon baseline models. Chen et al. high-frequency data was employed in [5] to implement LSTM-based feature learning, which revealed significantly better performance when compared against traditional architectures, including radial basis function (RBF) networks and extreme learning machines on Chinese equity markets, where performance scaling is favorable with respect to the volume of the training set. For example, Hammoudeh et al. [9] explored oil-renewable energy correlations through quantile-based causality frameworks, which located regime-dependent volatility spillovers.

2.2 Metaheuristic Hyperparameter Optimization

An increasing frequency of work applies nature-inspired optimization to neural network configuration. Swarm methods and evolutionary strategies such as PSO variants, genetic algorithms, and differential evolution have shown great promise in hyperparameter tuning of recurrent models [17]. More recently, quantum-inspired variants of PSO have gained interest: these methods realize stronger global search by batched probabilistic sampling over positions derived from the quantum wavefunction while avoiding stagnation through classical moment updates[18,15]. In a similar vein, chaotic maps have been employed in metaheuristic frameworks to substitute pseudo-random number generation, which results in population trajectories that are more ergodic and diverse [11].

2.3 Research Gaps

Despite this advancement, some limitations persist. First, most hyperparameter optimization studies that utilize quantum-inspired PSO or chaos-enhanced PSO alone do not seem to integrate both mechanisms into a unified framework. Second, systematic benchmarking of these hybridized approaches against a diverse range of gold-standard metaheuristics on unified test suites (like CEC-2022) is rather limited. Third, despite the aforementioned studies employing hyperparameter tuning using metaheuristics to predict financial time series in a variety of domains, they have been applied only to single markets or stocks and so provide indirect evidence of cross-sector generalizability. Fourth, computational complexity and training time critical for practical deployment are rarely reported alongside prediction accuracy. These gaps motivate the present work.

3 Proposed Model

3.1 Particle Swarm Optimization Foundation

PSO [3] maintains a swarm of N particles, each with position Z_{ij}^h and velocity V_{ij}^h in a P-dimensional search space. At each iteration h, velocities and positions are updated as:

$$V_{ij}^{h+1} = \omega \cdot V_{ij}^h + E_p \eta_p \left(Z_{\text{pbest}ij}^h - Z_{ij}^h \right) + E_g \eta_g \left(Z_{\text{gbest}i}^h - Z_{ij}^h \right), \tag{1}$$

$$Z_{ij}^{h+1} = Z_{ij}^h + V_{ij}^{h+1}, \quad i = 1, \dots, N, \quad j = 1, \dots, P, \tag{2}$$

where ω is the inertia weight, E_p and E_g are cognitive and social acceleration coefficients, and $\eta_p, \eta_g \in [0, 1]$ are uniform random values. Personal best Z_{pbest} and global best Z_{gbest} update as:

$$Z_{\text{pbest}ij}^{h+1} = \begin{cases} Z_i^{h+1}, & f(Z_i^{h+1}) < f(Z_{\text{pbest}i}^h), \\ Z_{\text{pbest}i}^h, & \text{otherwise}, \end{cases} \tag{3}$$

$$Z_{\text{gbest}i}^h = \arg\min_i f(Z_{\text{pbest}ij}^h). \tag{4}$$

A key limitation of this scheme is that the velocity term can lead to rapid exploitation of local regions, increasing susceptibility to premature convergence on multimodal landscapes [10].

3.2 Quantum-Inspired Position Update

HQCPSO replaces velocity-based movement with a quantum mechanical framework. Each particle is modeled as a quantum entity confined in a δ-function potential well centered at the attractor point $H_{i,j}(h)$. In this model, the probability density of finding the particle at position Z is given by the squared modulus of the wave function:

$$|\psi(Z, h)|^2 \propto \exp\left(-\frac{2|Z - H_{i,j}(h)|}{K_{i,j}(h)/2} \right), \tag{5}$$

where $K_{i,j}(h)$ is the characteristic width (analogous to the de Broglie wavelength scale). Sampling from this distribution via the inverse cumulative distribution function (Monte Carlo inversion) yields the position update:

$$Z = H_{i,j}(h) \pm \frac{K_{i,j}(h)}{2} \ln\left(\frac{1}{U}\right), \quad U \sim \mathrm{Uniform}(0,1]. \tag{6}$$

This construction ensures that particles can, in principle, tunnel beyond local attraction basins with nonzero probability at every iteration, a property absent from classical PSO.

3.3 Chaotic Attractor Integration

Whereas the quantum update provides stochastic breadth, the attractor location $H_{i,j}(h)$ controls the direction of search. In standard QPSO, H is a convex combination of Z_{pbest} and Z_{gbest} weighted by independent uniform random variables $\lambda_1, \lambda_2 \in [0,1]$. In HQCPSO, the weight is replaced by a chaotic sequence chao(i) generated by one of ten classical chaotic maps (Chebyshev, Circle, Gauss/mouse, Iterator, Logistic, Piecewise, Sine, Singer, Sinusoidal, Tent; see Fig. 1):

$$H_{i,j}(h) = \frac{\mathrm{chao}(i)}{\lambda_1 + \lambda_2} Z^h_{\mathrm{pbest}ij} + \left(1 - \frac{\mathrm{chao}(i)}{\lambda_1 + \lambda_2}\right) Z^h_{\mathrm{gbest}i}. \tag{7}$$

The potential width $K_{i,j}(h)$ is computed from the mean attractor position $M_j(h)$:

$$K_{i,j}(h) = 2\eta \left| M_j(h) - Z_{i,j}(h) \right|, \quad M_j(h) = \frac{1}{N} \sum_{i=1}^{N} H_{i,j}(h). \tag{8}$$

Substituting Eqs. (7) and (8) into Eq. (6) gives the complete HQCPSO position update:

$$Z_{i,j}(h+1) = H_{i,j}(h) \pm \eta \left| M_j(h) - Z_{i,j}(h) \right| \ln\left(\frac{1}{U}\right). \tag{9}$$

The contraction–expansion coefficient η decreases linearly from 1.0 to 0.5 over the course of optimization, gradually narrowing the quantum potential well and shifting the balance from exploration toward exploitation.

3.4 Algorithm and Complexity

Algorithm 1 summarizes the complete HQCPSO procedure. The computational complexity per iteration is $\mathcal{O}(N \cdot P)$, matching standard PSO, since quantum and chaotic operations are $\mathcal{O}(1)$ per particle per dimension. Over $Maxit$ iterations, total complexity is $\mathcal{O}(N \cdot P \cdot Maxit)$, comparable to all baseline optimizers evaluated in this study.

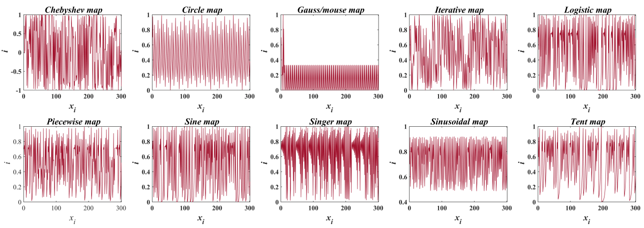

Fig. 1: Ten chaotic mapping functions used in HQCPSO. Each map generates a deterministic but ergodic sequence in $[0, 1]$, replacing the uniform random weights in standard QPSO to improve population diversity.

Algorithm 1 HQCPSO Algorithm

Require: Population size N, dimension P, max iterations $Maxit$, chaotic map selection
Ensure: Best solution Z_{gbest}

Initialize positions $Z^0_{i,j}$ randomly in $[\text{lb}_j, \text{ub}_j]$; evaluate fitness $f(Z^0_i)$
Set $Z_{\text{pbest}i} \leftarrow Z^0_i$; determine $Z_{\text{gbest}} \leftarrow \arg\min_i f(Z_{\text{pbest}i})$
for $h = 1$ to $Maxit$ **do**
  Linearly decrease $\eta : 1.0 \to 0.5$ ▷ Contraction–expansion schedule
  Generate chaotic sequence $\{\text{chao}(i)\}_{i=1}^{N}$ via selected map
  Compute mean attractor $M_j(h)$ via Eq. (8) for all j
  for $i = 1$ to N **do**
    for $j = 1$ to P **do**
      Draw $\lambda_1, \lambda_2 \sim \text{Uniform}(0, 1)$; draw $U \sim \text{Uniform}(0, 1)$
      Compute $H_{i,j}(h)$ via Eq. (7)
      Draw sign $s \in \{-1, +1\}$ with equal probability
      $Z_{i,j}(h + 1) \leftarrow H_{i,j}(h) + s\,\eta\,|M_j(h) - Z_{i,j}(h)|\,\ln(1/U)$
    end for
    Clip $Z_i(h + 1)$ to $[\text{lb}, \text{ub}]$; evaluate $f(Z_i(h + 1))$
  end for
  Update $Z_{\text{pbest}i}$ and Z_{gbest} if improved
end for

4 Experimental Setup

4.1 Benchmark Configuration

HQCPSO was tested on the CEC-2022 benchmark suite [1] that includes a total of 12 functions across unimodal (F1–F2), basic multimodal (F3–F7), hybrid (F8–F10), and composition (F11–F12) landscapes. For all algorithms, a swarm with $N = 30$ particles, $Maxit = 1000$ iterations, and 30 independent runs per function was used to ensure statistical reliability. Table 1 lists parameter settings.

Table 1: Algorithm parameter configurations

Algorithm	Parameters
PSO [3]	$c_1 = c_2 = 2$, $\omega \in [0.4, 0.9]$, $r_1, r_2 \in (0, 1)$
GOA [13]	$c \in [0.00004, 1]$
GWO [14]	$a : 2 \to 0$ (linear)
MO [19]	$P = 0.4$
HQCPSO (proposed)	$\omega \in [0.4, 0.9]$, $\lambda_1, \lambda_2 \in (0, 1)$, $\eta : 1.0 \to 0.5$

Experiments were conducted in MATLAB 2025b on an Intel Xeon W7-2495X processor (2.50 GHz, 18 cores) with 128 GB DDR4 RAM and NVIDIA RTX A4000 GPU acceleration.

4.2 Stock Market Dataset

Data were sourced from the National Stock Exchange (NSE) of India via the NSE official data portal , covering the five-year period 13/12/2018–13/12/2023 (approximately 1250 trading days per equity). Four indices were selected to represent diverse market sectors and volatility profiles: Reliance Industries (conglomerate/diversified), Adani Green Energy (renewable energy), Nestle India Limited (consumer goods/FMCG), and Coal India Limited (mining/energy). Each record includes opening price, daily high, daily low, closing price, volume-weighted average price, traded quantity, trade count, and deliverable volume. The dataset was partitioned 80% for training and 20% for testing (no overlap), with min-max normalization [2] applied feature-wise to the training set and subsequently applied to the test set using training-set statistics:

$$\xi_{\text{norm}} = \frac{\xi - \xi_{\min}}{\xi_{\max} - \xi_{\min}}. \tag{10}$$

4.3 LSTM Architecture and Hyperparameter Space

The LSTM-HQCPSO architecture (Table 2) comprises a fixed 20-neuron input layer, three stacked LSTM layers whose sizes Z_0, Z_2, Z_4 are determined by the optimizer, a dense layer of Z_5 neurons, a dropout layer with rate Z_6, and a single-neuron output. Layers are dynamically removed when the optimizer sets the corresponding flag (Z_1 or Z_3) to zero, enabling variable-depth architectures.

Table 2: LSTM-HQCPSO architecture specification

Stage	Layer	Configuration
1	Input	20 neurons (fixed)
2	LSTM-I	Z_0 neurons (optimized)
3	LSTM-II	Z_2 neurons (optimized; removed if $Z_1 = 0$)
4	LSTM-III	Z_4 neurons (optimized; removed if $Z_3 = 0$)
5	Dense	Z_5 neurons (optimized)
6	Dropout	Rate Z_6 (optimized)
7	Output	1 neuron (fixed)

4.4 Evaluation Metrics

Prediction performance is assessed using four complementary metrics:

$$\mathrm{MSE} = \frac{1}{N}\sum_{i=1}^{N}(\hat{\eta}_i - \eta_i)^2, \quad \mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}|\hat{\eta}_i - \eta_i|, \tag{11}$$

$$\mathrm{MAPE} = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{\eta_i - \hat{\eta}_i}{\eta_i}\right| \times 100, \quad R^2 = 1 - \frac{\sum_{i=1}^{N}(\eta_i - \hat{\eta}_i)^2}{\sum_{i=1}^{N}(\eta_i - \bar{\eta})^2}. \tag{12}$$

MAE measures average absolute deviation and is robust to outliers; MSE penalizes large errors quadratically, which is critical for risk-sensitive financial applications; MAPE provides a scale-free percentage error (interpreted cautiously when normalized prices approach zero); R^2 quantifies the proportion of price variance explained by the model.

5 Results and Analysis

5.1 Benchmark Performance

Table 3 presents mean and standard deviation results across 30 runs for all five algorithms. HQCPSO achieved rank-1 performance on 8 of 12 functions (F2, F3, F6, F7, F8, F9, F10, F11), with an overall mean rank of 1.5, substantially outperforming PSO (3.67), GOA (2.83), GWO (2.67), and MO (4.33).

Convergence analysis. Figure 2 shows convergence trajectories for all 12 functions. On multimodal functions F6 and F11, HQCPSO reaches near-optimal solutions within 200 iterations while competing algorithms plateau at suboptimal values. On F7, HQCPSO escapes a local optimum around iteration 100 where PSO and MO stagnate achieving 1.4% and 2.6% better final solutions respectively. On highly multimodal F10–F12, HQCPSO exhibits a three-phase trajectory: aggressive exploration (iterations 0–100), transitional balancing (100–300), and intensive exploitation (300–1000), consistent with its linearly decreasing η schedule.

Table 3: CEC-2022 benchmark results (mean ± std over 30 independent runs). Bold indicates best result per function.

Fn	Metric	HQCPSO	PSO	GOA	GWO	MO
F1	AVG	3.02e+02	1.74e+03	**3.00e+02**	1.73e+03	8.00e+03
	std	7.75e+00	1.47e+03	1.14e+00	1.53e+03	2.95e+03
F2	AVG	**4.07e+02**	4.26e+02	4.10e+02	4.28e+02	4.79e+02
	std	1.51e+01	2.44e+01	1.76e+01	2.11e+01	5.70e+01
F3	AVG	**6.01e+02**	6.18e+02	6.08e+02	6.02e+02	6.41e+02
	std	2.76e+00	8.61e+00	1.08e+01	1.88e+00	1.08e+01
F4	AVG	8.15e+02	8.29e+02	8.26e+02	**8.13e+02**	8.34e+02
	std	5.11e+00	1.21e+01	1.41e+01	5.12e+00	8.53e+00
F5	AVG	9.26e+02	1.22e+03	**9.02e+02**	9.13e+02	1.30e+03
	std	2.27e+01	3.22e+02	2.69e+00	2.27e+01	1.87e+02
F6	AVG	**2.17e+03**	5.07e+03	3.68e+03	5.75e+03	3.59e+03
	std	5.90e+02	2.67e+03	1.70e+03	2.42e+03	1.88e+03
F7	AVG	**2.03e+03**	2.06e+03	2.07e+03	2.03e+03	2.09e+03
	std	9.07e+00	2.92e+01	4.57e+01	1.39e+01	2.93e+01
F8	AVG	**2.22e+03**	2.24e+03	2.29e+03	2.23e+03	2.24e+03
	std	2.17e+00	3.65e+01	5.71e+01	2.29e+01	2.32e+01
F9	AVG	**2.53e+03**	2.55e+03	2.55e+03	2.58e+03	2.66e+03
	std	2.97e-12	4.48e+01	5.24e+01	4.66e+01	4.36e+01
F10	AVG	**2.53e+03**	2.61e+03	2.71e+03	2.56e+03	2.54e+03
	std	5.26e+01	1.41e+02	3.08e+02	5.74e+01	6.33e+01
F11	AVG	**2.66e+03**	2.87e+03	2.79e+03	2.79e+03	2.89e+03
	std	1.03e+02	1.84e+02	1.83e+02	1.71e+02	2.72e+02
F12	AVG	2.87e+03	2.88e+03	**2.86e+03**	2.87e+03	2.88e+03
	std	1.42e+00	2.27e+01	1.75e+00	3.72e+00	1.04e+01
Mean Rank		**1.5**	3.67	2.83	2.67	4.33

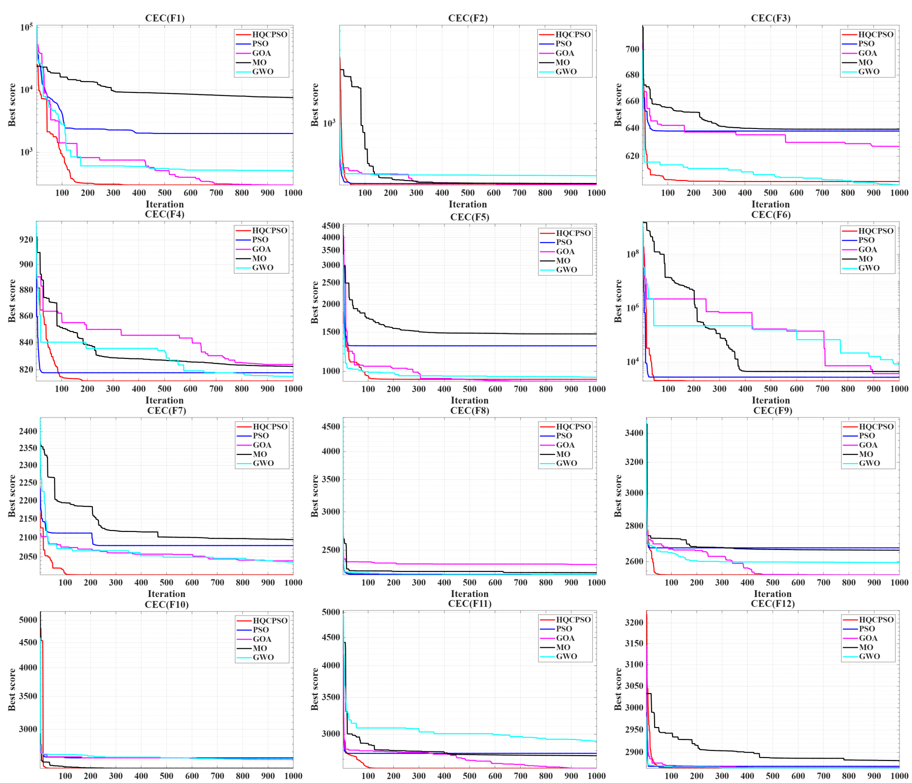

Fig. 2: Convergence curves for CEC-2022 benchmark functions F1–F12. Each curve shows the best fitness value averaged over 30 independent runs.

Exploration–exploitation dynamics. Figure 3 quantifies the balance between global search diversity and local intensification throughout optimization. Early iterations (0–50) exhibit 60–100% exploration, reflecting the quantum position update's probabilistic spread across the search space. As η decreases, exploitation dominance rises to 95–99% by iteration 200 on most functions. Function F6 maintains 3–5% residual exploration even after 800 iterations, attributed to chaotic perturbations that prevent full lock-in on the many local optima present in its landscape. Unimodal F2 transitions to 98% exploitation by iteration 80, reflecting the algorithm's adaptive calibration to simpler fitness topologies.

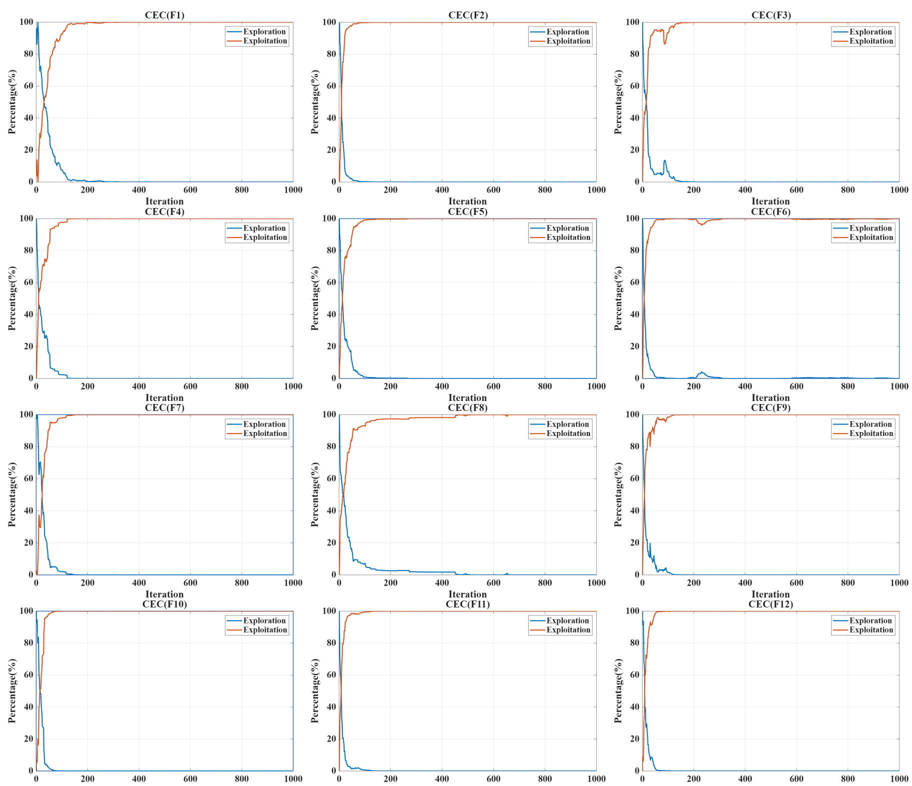

Fig. 3: Exploration–exploitation dynamics of HQCPSO across CEC-2022 functions F1–F12.

Statistical significance. Wilcoxon rank-sum tests (Table 4) confirm that HQCPSO's advantages are statistically significant against PSO ($p = 0.002104$), GWO ($p = 0.0425$), and MO ($p = 0.002264$). The GOA comparison yields $p = 0.06113$, marginally exceeding the 0.05 threshold; however, HQCPSO achieves superior mean rank (1.5 vs. 2.83) and wins on 8 versus GOA's 3 functions, indicating that GOA's strength is concentrated on specific function types (F1, F5, F12) while HQCPSO offers broader applicability.

Table 4: Wilcoxon rank-sum test results (HQCPSO vs. each baseline).

Statistic	**PSO**	**GOA**	**GWO**	**MO**
P-value	0.002104^{**}	0.06113	0.0425^{*}	0.002264^{**}
Z-statistic	-3.0751	-1.8725	-2.0286	-3.0532

$^{**}p < 0.01$; $^{*}p < 0.05$

Figure 4 provides a visual summary of mean ranks, Z-statistics, and p-values across all comparisons. Negative Z-statistics indicate HQCPSO wins.

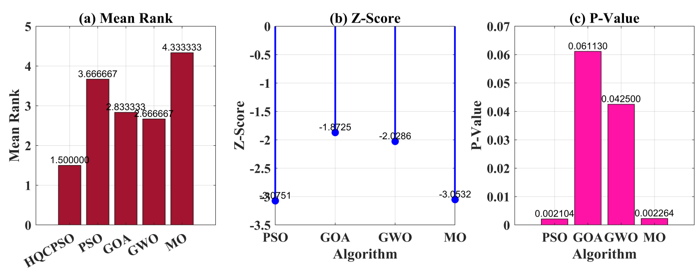

Fig. 4: Statistical comparison: (a) mean ranks, (b) Z-statistics, and (c) p-values of Wilcoxon rank-sum tests.

5.2 Stock Market Forecasting Results

Table 5 reports LSTM prediction performance across four NSE India equity indices under all five hyperparameter optimizers.

MAE and MSE. HQCPSO-optimized LSTM achieves the lowest MAE on all four datasets. For Reliance, MAE = 0.000664 represents a 27.7% improvement over PSO (0.000919), 33.3% over GOA (0.000995), 17.7% over GWO (0.000807), and 13.2% over MO (0.000765). The Nestle India dataset shows the most pronounced margin: HQCPSO's MAE of 0.001293 is 32.0% below PSO and 48.5% below GOA. MSE improvements are similarly consistent: 33.0% reduction (Coal India), 49.6% reduction (Nestle India), 16.2% (Adani Green), and 31.9% (Reliance), all relative to PSO. The quadratic penalization in MSE makes these improvements particularly meaningful for risk-sensitive applications where large prediction outliers carry disproportionate consequences.

MAPE. MAPE results require careful interpretation due to sensitivity to near-zero normalized price values. HQCPSO attains outstanding relative accuracy

Table 5: Stock price prediction performance comparison.

Dataset	Model	MAE	MSE	MAPE (%)	R^2
Coal India	HQCPSO	**0.000608**	**0.013536**	26.753	**0.940510**
	PSO	0.001006	0.020212	36.247	0.901570
	GOA	0.000622	0.014207	29.866	0.939140
	GWO	0.000681	0.014643	**18.931**	0.933290
	MO	0.000696	0.014656	24.963	0.931870
Nestle India	HQCPSO	**0.001293**	**0.015082**	**8.737**	**0.924880**
	PSO	0.001902	0.029916	16.802	0.889520
	GOA	0.002512	0.037035	23.229	0.854100
	GWO	0.001397	0.017955	11.374	0.918850
	MO	0.001482	0.018840	11.543	0.913900
Adani Green	HQCPSO	**0.000910**	**0.010794**	**217.94**	**0.922630**
	PSO	0.000967	0.012881	358.15	0.917780
	GOA	0.001157	0.019836	607.51	0.901660
	GWO	0.001027	0.016172	421.43	0.912730
	MO	0.001544	0.027786	775.35	0.868780
Reliance	HQCPSO	**0.000664**	**0.010997**	**2.715**	**0.916540**
	PSO	0.000919	0.016157	4.007	0.884480
	GOA	0.000995	0.020381	5.352	0.874970
	GWO	0.000807	0.012959	3.176	0.898620
	MO	0.000765	0.013731	3.533	0.903810

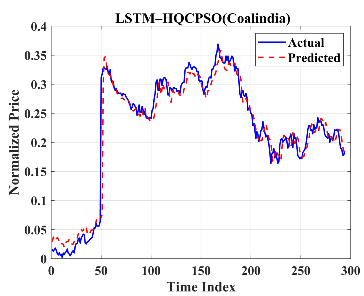

(a) Coal India Limited

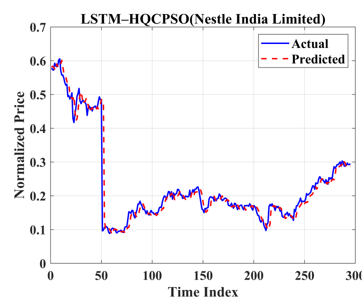

(b) Nestle India Limited

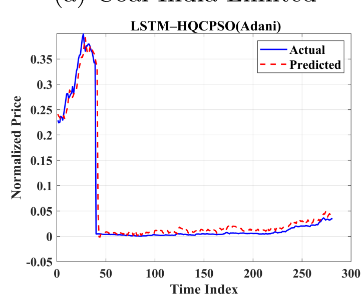

(c) Adani Green Energy

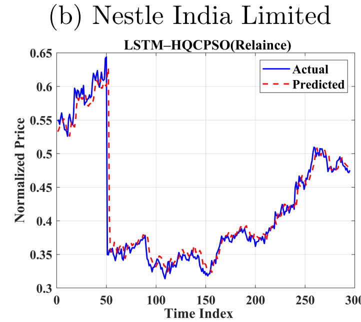

(d) Reliance Industries

Fig. 5: LSTM-HQCPSO predictions (red dashed) vs. actual normalized prices (blue solid) on the 20% test set.

for stable large-cap stocks (Reliance: 2.715%, Nestle India: 8.737%). But Coal India offers an exception to the rule, with GWO (18.931%) yielding better performance than HQCPSO (26.753%), but HQCPSO still leads on MAE and MSE metrics. This anomaly occurs because MAPE's denominator also becomes very small when normalized prices approach its lower limit, which seems to magnify percentage-based error, seemingly regardless of predictive quality. For Adani Green, extreme MAPE values (HQCPSO: 217.94%) in this class of stocks are attributable to the dramatic early price spike seen between 2018 and 2019 and the near-zero normalized values that followed, meaning MAPE is less informative than MAE and MSE as a primary metric for high beta equities.

R^2 and visual fit. HQCPSO achieves the highest R^2 across all datasets: Coal India (0.941), Nestle India (0.925), Adani Green (0.923), and Reliance (0.917), indicating that the optimized LSTM networks explain 92–94% of out-of-sample price variance.

Figure 5 shows time-series overlays of actual and predicted prices. In Fig. 5a, HQCPSO tracks Coal India's cyclical movements, including the sharp uptrend (indices 40–70, normalized value 0.05→0.33) and the subsequent downtrend (indices 180–240, 0.37→0.17). In Fig. 5b, Nestle India's consolidation phase (indices 60–240) and late uptrend (indices 240–300) are accurately captured. In Fig. 5c, HQCPSO handles Adani Green's early volatility spike and then reliably tracks the extended low-volatility phase (indices 50–250). In Fig. 5d, Reliance's U-shaped trajectory initial decline (indices 0–100), trough consolidation (100–200), and recovery (200–300) is reproduced with consistently tight alignment.

Computational overhead. The per-iteration cost of HQCPSO is $\mathcal{O}(N \cdot P)$, identical to standard PSO; the chaotic map evaluation and quantum sampling operations add negligible constant overhead. Average wall-clock time per optimizer run on the NSE datasets (30 optimizer iterations $\times$ LSTM training) was approximately 12–15 minutes per stock on the hardware described, with HQCPSO within 5% of PSO runtime. This validates that these performance gains come with no meaningful computational end cost compared to the baseline.

6 Conclusion

This paper proposes HQCPSO, a hybrid optimizer that uses δ potential-well quantum position updates and chaotic attractor sequences to avoid premature convergence present in classical PSO. Specifically, the quantum update obtains principled probabilistic exploration in the absence of velocity parameters and chaotic sequences to supersede pseudo-random weights with ergodic, diversity-preserving iterates. Benchmark validation on CEC-2022 supports rank-1 (mean rank 1.5) performance vs. PSO, GWO, and MO with statistical significance that $p < 0.05$ has been verified. Further experimentation involves LSTM hyperparameter tuning for NSE India equity prediction, where HQCPSO outperforms all other algorithms consistently across MAE and MSE and R^2 across four distinct market areas. The computational overhead over traditional PSO is minimal.

Future work may involve designing a multi-objective HQCPSO to trade off between prediction accuracy and computational cost to be used for real-time trading. In this paper HQCPSO was used only to optimize a single LSTM model, so its usage on ensemble LSTM models may help combine different results and provide better robustness and lower-variance predictions. The presented method would benefit from a more extensive baseline comparison versus modern optimization methods, like Bayesian optimization and CMA-ES. The generalizability of the model could be verified by testing it on other stock markets such as NYSE and HKEX. Furthermore, adapting the framework to intraday and high-frequency trading data would need additional consideration of more intricate market dynamics. Future research will also compare the proposed model with traditional econometric methods and new deep learning architectures like Transformers.

Acknowledgments. No external funding was received for this research. The authors also acknowledge the National Institute of Technology Agartala for providing computational resources.

Disclosure of Interests. The authors declare no competing interests relevant to the content of this article.

References

1. Ahrari, A., Elsayed, S., Sarker, R., Essam, D., Coello, C.A.C.: Problem definition and evaluation criteria for the cec'2022 competition on dynamic multimodal optimization. In: Proceedings of the IEEE World Congress on Computational Intelligence (IEEE WCCI 2022), Padua, Italy. pp. 18–23 (2022)
2. Ali, P.J.M., Faraj, R.H., Koya, E., Ali, P.J.M., Faraj, R.H.: Data normalization and standardization: a technical report. Mach Learn Tech Rep **1**(1), 1–6 (2014)
3. Bratton, D., Kennedy, J.: Defining a standard for particle swarm optimization. In: 2007 IEEE swarm intelligence symposium. pp. 120–127. IEEE (2007)
4. Chen, H., Zhou, Y., Cao, Q.: A hybrid learning framework integrating chaotic niche alpha evolution for student academic performance prediction. Scientific Reports (2026)
5. Chen, L., Qiao, Z., Wang, M., Wang, C., Du, R., Stanley, H.E.: Which artificial intelligence algorithm better predicts the chinese stock market? IEEE Access **6**, 48625–48633 (2018)
6. Chen, M.Y., Liao, C.H., Hsieh, R.P.: Modeling public mood and emotion: Stock market trend prediction with anticipatory computing approach. Computers in Human Behavior **101**, 402–408 (2019)
7. DiPietro, R., Hager, G.D.: Deep learning: Rnns and lstm. In: Handbook of medical image computing and computer assisted intervention, pp. 503–519. Elsevier (2020)
8. Gülmez, B.: Ga-attention-fuzzy-stock-net: An optimized neuro-fuzzy system for stock market price prediction with genetic algorithm and attention mechanism. Heliyon **11**(3) (2025)
9. Hammoudeh, S., Mokni, K., Ben-Salha, O., Ajmi, A.N.: Distributional predictability between oil prices and renewable energy stocks: is there a role for the covid-19 pandemic? Energy Economics **103**, 105512 (2021)

10. Jiang, Y., Hu, T., Huang, C., Wu, X.: An improved particle swarm optimization algorithm. Applied Mathematics and Computation **193**(1), 231–239 (2007)
11. Limane, A., Zitouni, F., Harous, S., Lakbichi, R., Ferhat, A., Almazyad, A.S., Jangir, P., Mohamed, A.W.: Chaos-enhanced metaheuristics: classification, comparison, and convergence analysis. Complex & Intelligent Systems **11**(3), 177 (2025)
12. Malik, S., Patro, S.G.K., Mahanty, C., Lasisi, A., Naveed, Q.N., Buradi, A., Emma, A.F., Kumar, S., Mubarakali, A.: Quantum-inspired gravitationally guided particle swarm optimization for feature selection and classification. Scientific Reports **15**(1), 34155 (2025)
13. Meraihi, Y., Gabis, A.B., Mirjalili, S., Ramdane-Cherif, A.: Grasshopper optimization algorithm: theory, variants, and applications. Ieee Access **9**, 50001–50024 (2021)
14. Mirjalili, S., Mirjalili, S.M., Lewis, A.: Grey wolf optimizer. Advances in engineering software **69**, 46–61 (2014)
15. Nagaraju, J., Vijaya, J.: Boost customer churn prediction in the insurance industry using meta-heuristic models. International Journal of Information Technology (2022)
16. Rezaei, H., Faaljou, H., Mansourfar, G.: Stock price prediction using deep learning and frequency decomposition. Expert Systems with Applications **169**, 114332 (2021)
17. Vijaya, J., Rao, M.: Heart disease prediction using clustered particle swarm optimization techniques. In: 2022 IEEE 6th Conference on Information and Communication Technology (CICT). IEEE (2022)
18. Vijaya, J., Sivasankar, E.: An efficient system for customer churn prediction through particle swarm optimization based feature selection model with simulated annealing. Cluster Computing (2017)
19. Yuan, Y., Chong, G., Ren, J., Zhao, W., Li, Y., Wang, Z., Mirjalili, S.: Musk ox optimizer (mo): a novel optimization algorithm and its application. Cluster Computing **28**(16), 1041 (2025)
20. Zhou, J., Lin, H., Chen, Y.: A high-efficiency and interpretable framework for microparameter calibration via pso-optimized machine learning and shap analysis. Journal of Materials Engineering and Performance pp. 1–22 (2026)