

MACPGANA Under Data Scarcity: Single-Commodity Validation of a GAN–Autoencoder–Q-Learning Framework for Agricultural Price Forecasting

Rubi Kambo^{1,2}, Jyothi Pillai³, Sunita Soni⁴

¹ Department of Computer Science and Engineering, Bhilai Institute of Technology, Durg, Chhattisgarh, India.

² Department of Computer Science and Engineering, Rungta College of Engineering and Technology, Rungta International Skills University, Bhilai, Chhattisgarh, India.

Email: rubi.kambo7@gmail.com

³ Department of Information Technology, Bhilai Institute of Technology, Durg, Chhattisgarh, India.

Email: jyothi.pillai@bitdurg.ac.in

⁴ Department of Computer Science and Engineering, Bhilai Institute of Technology, Durg, Chhattisgarh, India.

Email: sunita.soni@bitdurg.ac.in

Corresponding Author: Rubi Kambo

Abstract: Accurate short-term forecasting of agricultural commodity prices is critical for protecting farmer incomes, guiding traders' inventory decisions, and informing policy interventions. However, most deep learning forecasting frameworks are validated only on large, multi-commodity datasets, leaving their behaviour under the data-scarce, single-commodity conditions typical of regional agricultural markets largely unexamined. This study addresses that gap by evaluating MACPGANA — a hybrid framework integrating multimodal feature extraction, Generative Adversarial Network (GAN)-based feature selection, autoencoder-based price prediction, and Q-Learning-driven adaptive optimisation against three established baseline architectures (Multilayer Perceptron, Convolutional Neural Network, and Long Short-Term Memory) under a constrained single-commodity setting. All four models were trained and evaluated under identical conditions on 365 days of daily wholesale tomato price data from the Raipur Agricultural Produce Market Committee, Chhattisgarh, India (2018), using a chronological train–test split. MACPGANA achieved a Mean Absolute Percentage Error of 1.66% and a normalised RMSE of 0.0128 (₹29.38/quintal), corresponding to MSE reductions of 93.0–97.0% relative to the baselines. Under a $\pm 10\%$ tolerance-band classification criterion, MACPGANA achieved 94.59% accuracy and an F1-score of 97.22%, against 29.73–43.33% accuracy for the baselines, with directional accuracy of 91.67% versus 0–8.33%. These results show that a framework engineered for large-scale, multimodal price prediction retains substantial predictive advantage when applied to a single commodity with limited historical data, supporting its applicability to resource-constrained regional markets.

Keywords: Agricultural commodity price forecasting; Generative Adversarial Networks; Autoencoder; Q-Learning; Deep learning; LSTM; Single-commodity validation; Data-scarce time series

1. Introduction

Agricultural commodity price forecasting plays a critical role in supporting decisions made by farmers, traders, and policymakers across emerging economies. In India, agriculture and allied sectors contribute approximately 18% of national Gross Value Added and engage over half of the total workforce (54.6%, per the 2011 Census) [1], and price volatility in perishable commodities such as tomato which can swing nearly nine-fold within a single growing season at a regional market translates directly into income instability for smallholder producers [2,3].



Deep learning architectures have increasingly displaced classical statistical forecasting tools such as ARIMA and its seasonal variants, which assume linearity and stationarity that agricultural price series routinely violate [4]. Feedforward networks, convolutional architectures, and recurrent models with gated memory cells have each been shown to capture non-linear and temporally structured price dynamics more effectively than statistical baselines [5,6]. More recently, hybrid frameworks combining generative adversarial feature selection, autoencoder-based representation learning, and reinforcement-learning-driven adaptation have been proposed to address the multimodal, non-stationary character of commodity markets at scale. One such framework, MACPGANA, was validated on a 100,000-record, seven-commodity international dataset and reported substantial accuracy gains over recurrent, feedforward, and regression baselines [7].

A framework validated only at large scale, however, leaves an important practical question unanswered: does its advantage persist under the single-commodity, short-history conditions that characterise most regional agricultural markets in practice? Many Agricultural Produce Market Committees in India have recorded electronic price data for only a decade or two, and a given commodity-region combination may offer only a single growing season of usable observations [8]. Existing comparative studies of deep learning architectures for agricultural price forecasting [9,10] rarely test a multimodal hybrid architecture against simpler baselines under such constrained conditions within a single, controlled experimental design, leaving practitioners without clear evidence on whether architectural complexity remains justified once data is scarce.

The study addresses that question directly. I evaluate MACPGANA against three widely used baseline architectures — Multilayer Perceptron (MLP), Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM) — on 365 days of daily wholesale tomato price data from the Raipur Agricultural Produce Market Committee, Chhattisgarh, India, training and testing all four models under identical conditions. The objective is to determine whether MACPGANA retains its forecasting advantage when applied outside the large-scale, multimodal setting for which it was originally designed.

The key contributions of this paper are as follows.

1. A controlled, single-commodity comparative evaluation of MACPGANA against MLP, CNN, and LSTM under identical data, preprocessing, and training conditions, closing a validation gap left open by MACPGANA's original large-scale evaluation.
2. A dual evaluation framework combining regression accuracy (RMSE, MAE, MAPE, R^2) with a tolerance-band classification reading and a directional-accuracy reading, both explicitly derived from the same continuous price output.
3. Evidence that a hybrid GAN–autoencoder–Q-Learning architecture retains the large majority of its predictive advantage even when trained on as few as 250 daily observations, with practical implications for deployment in data-scarce regional markets.

The remainder of this paper is organised as follows: Section 2 reviews related work on statistical, machine learning, and deep learning approaches to agricultural price forecasting; Section 3 states the research gap explicitly; Section 4 describes the methodology and dataset; Section 5 presents experimental results; Section 6 discusses findings in relation to prior work; and Section 7 concludes with limitations and future directions.

2. Literature Review

2.1 Statistical and Econometric Approaches

Classical time-series methods dominated agricultural price forecasting for several decades before deep learning approaches gained traction. ARIMA and its seasonal extension remain attractive for their interpretability and modest data requirements [4]. Dharavath and Khosla [11] applied a seasonal ARIMA model to forecast fruit and vegetable prices in an Indian metropolitan market, while Bakhtadze, Maximov, and Maximova [12] developed regression-based local wheat price prediction models for a constrained regional context. Hegde, Hulipalled, and Simha [13] conducted an early comparative study of statistical and machine-learning techniques for agricultural commodity price prediction, and Sun et al. [14] subsequently provided a comprehensive review consolidating statistical, machine-learning, and hybrid forecasting methods across diverse commodity classes. Addressing the inflationary dimension of price formation specifically, Dutta et al. [15] proposed a neuro-fuzzy system for predicting agricultural inflation, while Zhang, Chen, Liwen, and Xia [16] introduced a model-selection framework incorporating time-series features and

forecast horizons to guide technique choice. Liu and Lv [17] demonstrated that improved data-mining preprocessing could meaningfully enhance the accuracy of commodity price evaluation even within a largely statistical pipeline.

2.2 Machine Learning and Early Neural Approaches

The limitations of linear statistical models — chiefly their inability to capture non-linear, regime-dependent price behaviour — motivated a shift toward machine learning. Rashid et al. [18] reviewed machine-learning approaches to crop-yield prediction with a particular emphasis on palm oil, surveying regression, ensemble, and early neural techniques that later informed price-forecasting pipelines. Hegde, Hulipalled, and Simha [19] extended their earlier statistical comparison by incorporating natural language processing alongside machine learning for agricultural price prediction. Sharma, Sharma, and Verma [20] examined machine-learning algorithms for crop yield prediction specific to Indian agricultural contexts, and Adhikari et al. [21] built a market-intelligence system for agricultural commodities using cardamom as a case study, integrating multiple data streams into a decision-support pipeline. Havaluddin et al. [22] applied a backpropagation neural network to agricultural product price prediction, an early demonstration that feedforward neural architectures could outperform purely statistical baselines in this domain. Urkude and Pandey [23] developed AgriSense, an IoT-based irrigation utility system whose embedded ANN component has subsequently served as a comparative benchmark in later multimodal forecasting studies.

2.3 Deep Learning Architectures: ANN, CNN, and LSTM

As deep learning matured, convolutional and recurrent architectures were increasingly applied to agricultural time series. Banerjee, Sinha, and Choudhury [24] demonstrated long- and short-term forecasting of horticultural produce prices using an LSTM network model, reporting clear gains over regression baselines. Joshi and Patel [25] proposed a CNN–Bidirectional LSTM hybrid for agricultural commodity price forecasting in the Gujarat region, combining convolutional feature extraction with bidirectional sequence modelling. Manogna and Mishra [26] applied deep-learning models to forecast spot prices of agricultural commodities across Indian markets, while Zhang and Na [27] proposed a fuzzy-information-granulation and support-vector-machine hybrid optimised via a mind evolutionary algorithm, illustrating the broader trend toward combining classical learning algorithms with intelligent preprocessing. Gupta et al. [28] developed a machine-learning and feature-selection-enabled framework for crop yield prediction, underscoring the importance of feature engineering quality as a determinant of downstream forecasting accuracy. Vohra, Pandey, and Khatri [29] built a decision-support system for agricultural commodity price prediction, and Mahto et al. [30] applied artificial neural networks for short-term forecasting of agricultural commodities specifically within the Indian market context, reporting that sustainable agricultural decision-making depends critically on horizon-appropriate model selection.

2.4 Volatility, Futures, and Macro-Financial Price Dynamics

A parallel body of work has examined agricultural commodity price dynamics through an econometric and financial-markets lens. Fang, Guan, Wu, and Heravi [31] proposed an optimal forecast-combination strategy based on ensemble empirical mode decomposition for agricultural commodity futures prices, and Kurumatani [32] evaluated recurrent neural networks for time-series forecasting of agricultural product prices, proposing a dedicated evaluation methodology for this setting. Van Oordt, Stork, and De Vries [33] studied extreme price risk in agricultural commodities, while Ouyang, Wei, and Wu [34] proposed a long- and short-term time-series network for agricultural commodity futures price prediction. Marfatia, Ji, and Luo [35] examined volatility spillovers between agricultural commodity futures and oil markets, and Crespo Cuaresma, Hlouskova, and Obersteiner [36] developed a comprehensive econometric approach to agricultural commodity price dynamics and their macro-level determinants. Xu and Zhang [37] forecast prices for eight major agricultural commodities including coffee, corn, cotton, and wheat using neural networks, and He et al. [38] projected the effects of China’s corn and biofuel policy shifts on international agricultural trade using a commodity market model, illustrating how policy and trade variables interact with price formation at scale.

2.5 Generative, Hybrid, and Reinforcement-Learning-Based Approaches

Most directly relevant to the present study, recent work has explored generative and adaptive architectures for price and time-series forecasting. Diqi, Hiswati, and Nur [39] proposed StockGAN, demonstrating that GAN-based architectures can outperform traditional methods for volatile financial price prediction, a finding with direct relevance to commodity markets despite the non-agricultural domain. Goel and Mishra [40] surveyed recent deep-learning algorithms used in smart farming, situating GAN- and autoencoder-based methods within a broader and rapidly evolving methodological landscape. Hajirahimi, Khashei, and Hamadani [41] proposed a principal-component-based

hybrid model for general time-series forecasting, addressing multi-collinearity issues that also affect high-dimensional agricultural feature sets. The reinforcement-learning foundation underlying adaptive optimisation in hybrid forecasting systems traces to Sutton and Barto’s foundational text [42], while the generative adversarial framework itself originates with Goodfellow et al. [43]. Building on this foundation, MACPGANA [7] integrated multimodal feature extraction, GAN-based feature selection, autoencoder-based prediction, and Q-Learning-driven adaptive optimisation into a single framework, reporting state-of-the-art accuracy on a 100,000-record, seven-commodity international dataset.

2.6 Synthesis

Across these five strands of literature, a consistent pattern emerges: statistical models offer interpretability but fail under non-linearity [4,11,12], early machine-learning and shallow neural methods improve accuracy but typically operate on a single data modality [18–23], deep architectures (CNN, LSTM) capture temporal structure more effectively but remain static post-deployment [24–30], econometric and financial-markets studies highlight volatility and macro determinants largely outside the scope of price-history-only models [31–38] and generative/hybrid architectures though promising have so far been validated predominantly at large scale, multimodal, or non-agricultural settings [7,39–43]. None of the reviewed studies provides a controlled, identical-conditions comparison of a GAN–autoencoder–Q-Learning hybrid against standard MLP, CNN, and LSTM baselines under the single-commodity, data-constrained conditions that define most regional agricultural markets the specific gap this paper addresses in Section 3.

3. Research Gap

Section 2 establishes that statistical, machine learning, and deep learning approaches to agricultural commodity price forecasting have each been investigated extensively, and that generative and hybrid architectures represent the current methodological frontier. A specific gap, however, remains unaddressed: no published study has evaluated a multimodal GAN–autoencoder hybrid framework against standard MLP, CNN, and LSTM architectures under a controlled, identical-conditions comparison restricted to a single commodity with limited historical observations. MACPGANA’s original validation [7] employed a 100,000-record, seven-commodity dataset; whether its architectural advantages multimodal feature integration, adversarial feature selection, and Q-Learning-driven adaptation persist when training data is reduced by roughly three orders of magnitude, to the 250-observation regime typical of a single regional market and growing season, has not previously been established.

This represents a practically significant validation gap: practitioners deploying deep learning systems in data-scarce regional markets cannot determine, from the existing literature, whether architectural sophistication remains justified once large-scale, multimodal data is unavailable. The present study closes this gap through a controlled, single-commodity, identical-conditions evaluation of MACPGANA against three standard baseline architectures.

4. Methodology

4.1 Dataset

This study uses daily modal (wholesale) tomato price records from the Raipur Agricultural Produce Market Committee (APMC), Chhattisgarh, India, for calendar year 2018, obtained from the National Horticulture Board’s public price-reporting portal. The Raipur APMC is the principal wholesale hub for perishable horticultural produce in central Chhattisgarh. Tomato was selected for its high intra-year price volatility and its economic significance to smallholder farmers in the region. Of 365 calendar-day records, approximately 235 corresponded to active trading days after excluding weekly market closures and gazetted public holidays. Modal prices ranged from ₹300 to ₹2,600 per quintal (100 kg), with a lower-volatility sub-period (₹1,000–1,200/quintal) in the post-monsoon October–December window forming the test partition; January–August data formed the training partition. Partitioning preserved chronological order throughout, with no random shuffling, to prevent information leakage from future observations into model training.

4.2 Data Preprocessing

All price series were normalised using min-max scaling computed exclusively from training-partition statistics, then mapped to a $[-1, 1]$ range for the MLP, CNN, and LSTM baselines. For MACPGANA, the same training-partition price range (₹2,300, i.e., ₹2,600 – ₹300) was used to compute normalised regression metrics (NRMSE, NMAE), enabling scale-independent comparison across all four models. For CNN and LSTM, sliding input windows of width 3 and 7 days respectively were used, consistent with window-size optimisation conducted in preliminary baseline

experiments. This windowing reduces the usable test set to 36 sequential predictions, the basis for all classification and directional-accuracy figures reported in Section 5.

4.3 Baseline Architectures

Three baseline architectures were implemented as comparators, each representing a distinct level of representational capability.

Multilayer Perceptron (MLP): A feedforward network with three hidden layers (128→64→32 neurons, ReLU activation, dropout $p = 0.2$ on the first two layers) and a single linear output neuron, trained with the Adam optimiser (learning rate 0.001) and mean squared error (MSE) loss.

Convolutional Neural Network (CNN): A five-layer one-dimensional convolutional architecture processing 3-day input windows: a Conv1D layer (64 filters, kernel size 2, ReLU) followed by max-pooling (pool size 2), a flatten operation, a dense layer (50 neurons, ReLU), and a linear output neuron, totalling 3,493 trainable parameters.

Long Short-Term Memory (LSTM): A recurrent architecture with a 7-day input window feeding a single LSTM layer of 50 units (tanh activation), followed by a dropout layer (rate 0.2) and a linear output neuron, trained with the Adam optimiser and early stopping (patience = 30 epochs, restoring best validation weights).

4.4 Proposed MACPGANA Architecture

MACPGANA (Multimodal Agriculture Commodity Price prediction via Generative Adversarial Networks and Autoencoders) integrates four sequential stages, as illustrated in Fig. 1. First, a multimodal feature-extraction layer constructs a high-density feature vector (HDFV) from the raw price series, comprising technical indicators (simple, exponential, and adaptive moving averages; Aroon; Average True Range; Bollinger Bands; Commodity Channel Index; MACD; On-Balance Volume; Psychological Line; Rank Correlation Index), frequency-domain transforms (discrete Fourier and discrete cosine transforms), and convolutional features, together yielding a super feature vector (SFV) of more than 60 candidate predictors [7]. Second, a Generative Adversarial Network performs adaptive feature selection: a discriminator evaluates variance-maximising subsets of the SFV against ground-truth price movements, retaining approximately 36 high-predictive-value features for the single-commodity configuration evaluated in this study. Third, the selected features are passed through a symmetric autoencoder (encoder 64→32→16, with a 16-dimensional latent bottleneck and a mirrored decoder) that compresses the feature space while preserving predictive signal, from which the next-day price prediction is generated. Fourth, a Q-Learning module compares predicted price movement against ground truth across consecutive iterations and adaptively tunes the model’s quantisation-error term, providing continuous post-training optimisation without requiring full retraining.

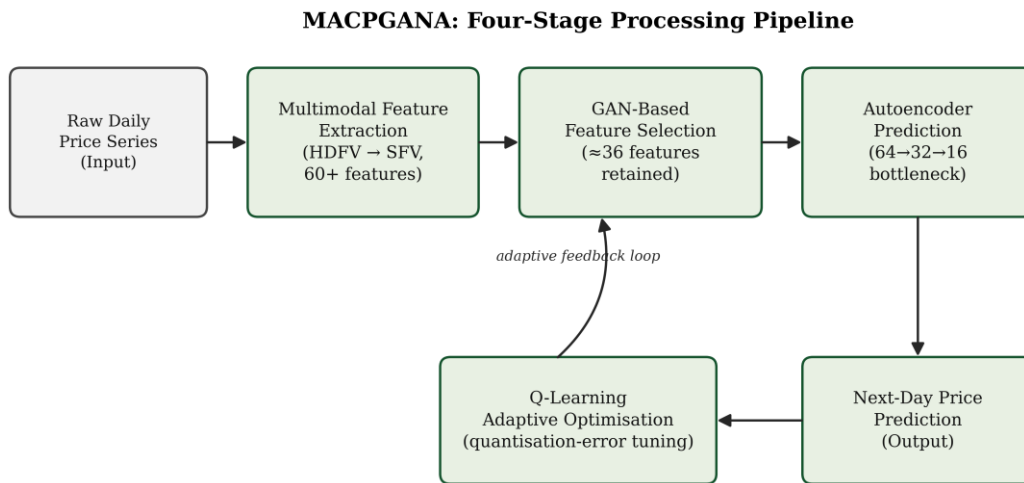


Fig. 1. MACPGANA four-stage processing pipeline: multimodal feature extraction, GAN-based feature selection, autoencoder-based prediction, and Q-Learning adaptive optimisation.

The complete procedure is formalised in Algorithm 1, which traces each of the four stages described above from raw price input to the adaptively corrected output prediction.

Algorithm 1. MACPGANA Single-Commodity Price Prediction

```

Input: Daily price series  $P = \{p_1, p_2, \dots, p_N\}$ ; GAN variance threshold  $\theta$ ;
       autoencoder bottleneck dimension  $d$ ; learning rate  $L_r$ 
Output: Predicted price  $AC\_out(t)$  for each test-partition day  $t$ 

// Stage 1: Multimodal Feature Extraction
1: for  $i = 1$  to  $N$  do
2:   Compute HDFV $i$  from raw price  $p_i$ 
3:   Compute technical indicators: SMA $i$ , EMA $i$ , AMA $i$ , Aroonup, $i$ ,
       Aroondown, $i$ , ATR $i$ , BB $i$ , CCI $i$ , MACD $i$ , OBV $i$ , PSL $i$ , RCI $i$ 
4:   Compute frequency-domain features:  $F_i$  (DFT), DCT $i$ 
5:   Compute convolutional feature: Convout, $i$ 
6:   Concatenate all features  $\rightarrow$  Super Feature Vector SFV $i$ 
7: end for

// Stage 2: GAN-Based Feature Selection
8: Initialise generator  $G$  and discriminator  $D$ 
9: for each training epoch do
10:  Compute discriminator loss  $L(SFV) = SFV\_g \cdot \log(SFV\_t) +$ 
        $(1 - SFV\_g) \cdot \log(1 - SFV\_t)$ 
11:  Compute max-variance objective  $L(Max)$  and
       min-variance objective  $L(Min)$ 
12:  Update  $G, D$  via gradient descent on gain function
        $G = \text{Min}[\text{Max}[\log(f\_t) + \log(1 - f\_g)]]$ 
13: end for
14: Compute GAN threshold  $GAN\_th$ 
15:  $SFV\_sel \leftarrow \{f \in SFV : \text{variance}(f) > GAN\_th\}$  //  $\approx 36$  features retained

// Stage 3: Autoencoder-Based Prediction
16: Encode  $SFV\_sel$  through encoder ( $64 \rightarrow 32 \rightarrow 16$ )
17: Decode latent representation through mirrored decoder
18: Compute  $AC\_out(t) = (1/N(SFV\_sel)) \Sigma \text{ARIMA}(t, SFV\_sel) \cdot \text{var}(SFV\_sel)$ 

// Stage 4: Q-Learning Adaptive Optimisation
19: for two consecutive iterations  $t, t+1$  do
20:  Compute prediction accuracy  $A(t) = [\text{Max}(P\_actual, P\_pred) -$ 
        $\text{Min}(P\_actual, P\_pred)] / \text{Max}(P\_actual, P\_pred)$ 
21:  Update quantisation error:
        $\epsilon(\text{New}) \leftarrow \epsilon(\text{Old}) + L_r \cdot [A(\text{Old}) + d \cdot \max(\epsilon) - \epsilon(\text{Old})]$ 
22:  Feed  $\epsilon(\text{New})$  back into Stage 2 (GAN threshold) and
       Stage 3 (ARIMA) for the next iteration

```

23: end for

24: return AC_out(t) for all t in the test partition

Note: Equation numbers correspond to the discriminator loss, GAN threshold, autoencoder prediction, and accuracy/Q-Learning update equations of the original MACPGANA formulation [7], here instantiated for the single-commodity configuration evaluated in this study.

4.5 Evaluation Metrics

Model performance was assessed using four regression metrics — Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) — together with the coefficient of determination (R^2). RMSE and MAE are additionally reported in normalised form, dividing by the ₹2,300 training-partition price range, to enable scale-independent comparison. Two complementary binary-label schemes were derived from the continuous price output purely for operational interpretation, and neither alters the underlying regression formulation: a tolerance-band scheme, under which a prediction is labelled correct if it falls within $\pm 10\%$ of the actual price (yielding Accuracy, Precision, Recall, and F1-score), and a directional scheme, under which a prediction is labelled correct if its sign matches the actual day-on-day price change (yielding Directional Accuracy). Mean inference latency was measured over 50 independent trials per model.

5. Experimental Setup

All four models were trained and evaluated under harmonised conditions to ensure a fair comparison: Adam optimiser (learning rate = 0.001), batch size 32, a maximum of 50 training epochs, and early stopping with patience of 10 epochs monitored on a held-out validation partition. CNN and LSTM used a 7-day input window; MACPGANA’s autoencoder operated on the GAN-selected feature subset derived from the same chronological partition described in Section 4.1. The chronological train–test split was preserved throughout, with identical preprocessing (min-max normalisation fitted on training data only) applied uniformly across all four architectures.

All models were implemented in Python 3.8 using TensorFlow 2.x and Keras for model construction and training, Scikit-learn (v0.24) for preprocessing and evaluation metrics, NumPy (v1.21) and Pandas (v1.3) for data manipulation, and Matplotlib (v3.4) / Seaborn (v0.11) for visualisation. Experiments were executed on a workstation running Ubuntu 20.04 LTS with an Intel Core i7-9700 CPU (8 cores, 3.0 GHz), 16 GB DDR4 RAM, and an NVIDIA GeForce RTX 2060 GPU (6 GB VRAM, CUDA 11.2). Random seeds were fixed globally (NumPy and TensorFlow seed = 42) at the start of every experimental run to ensure reproducibility.

6. Results

Table 1 presents consolidated regression and classification performance for all four models on the Raipur tomato test partition. MACPGANA achieves the lowest error on every regression metric and the highest score on every classification metric.

Table 1. Consolidated Regression, Classification, and Directional Performance

Model	MSE (norm. / ₹ ²)	RMS E (norm ./ ₹)	MAE (norm ./ ₹)	MAP E (%)	R^2	Accura cy (%)	Dir. Acc . (%)
MACPGA NA	0.000152 (805.62)	0.0128 (29.38)	0.0075 (17.16)	1.66	-0.5 8	94.59	91. 67
MLP	0.004573 (24,190.5 7)	0.0676 (155.5 3)	0.0607 (139.5 1)	13.93	-46. 31	29.73	8.3 3

CNN	0.005057 (26,752.5 3)	0.0711 (163.5 6)	0.0653 (150.1 2)	15.01	0.00	30.00	0.0 0
LSTM	0.002185 (11,558.7 1)	0.0467 (107.5 1)	0.0451 (103.7 1)	10.37	0.00	43.33	0.0 0

Note: All models use Adam optimiser; learning rate = 0.001, batch size 32, 50 epochs, early stopping patience = 10. CNN and LSTM use a 7-day input window. Classification criterion: prediction within $\pm 10\%$ of actual price.

On tolerance-based classification, MACPGANA achieves 94.59% accuracy, corresponding to 34 of 36 test predictions falling within the $\pm 10\%$ acceptance band — an improvement of 64.86 percentage points (pp) over MLP, 64.59 pp over CNN, and 51.26 pp over LSTM. Directional accuracy shows the most operationally significant gap: MACPGANA correctly predicts price-movement direction in 91.67% of day-to-day transitions, while MLP achieves only 8.33% and both CNN and LSTM achieve 0.00%, falling below the 50% random-guessing baseline.




Fig. 2. Comparison of (a) Root Mean Squared Error and (b) Mean Absolute Percentage Error across MACPGANA and the three baseline models.

Table 2 extends the evaluation to F1-score and inference latency, providing a complete assessment of operational deployment readiness.

Table 2. Classification Performance and Inference Latency

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Latency (ms)
MACPGANA	94.59	100.00	94.59	97.22	93.18
MLP	29.73	100.00	29.73	45.83	73.19
CNN	30.00	100.00	30.00	46.15	76.71
LSTM	43.33	100.00	43.33	60.47	77.08

Note: Latency is mean inference time over 50 independent trials. All four models achieve 100% precision under the $\pm 10\%$ criterion; recall is the discriminating factor, reflecting how completely each model captures genuinely acceptable predictions rather than defaulting to conservative near-mean outputs.

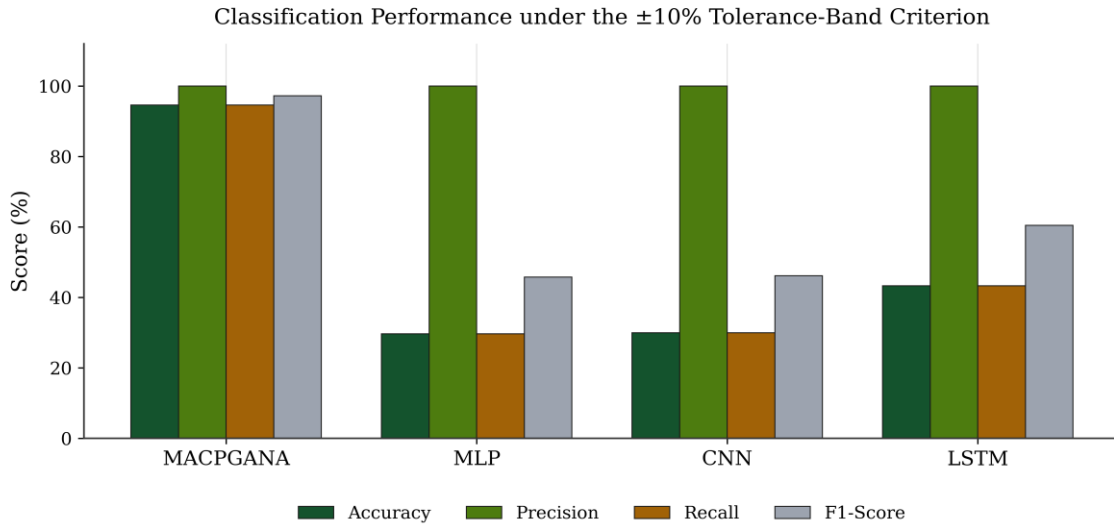


Fig. 3. Classification performance — Accuracy, Precision, Recall, and F1-Score — under the $\pm 10\%$ tolerance-band criterion.

Table 3 consolidates MACPGANA’s relative performance gains over each baseline across all primary metrics.

Table 3. MACPGANA Performance Improvement over Baseline Models

Performance Metric	vs MLP	vs CNN	vs LSTM
MSE Reduction	96.67%	96.99%	93.03%
RMSE Improvement	-81.1% (₹126.15)	-82.0% (₹134.18)	-72.7% (₹78.13)
MAE Improvement	-87.7% (₹122.35)	-88.6% (₹132.96)	-83.5% (₹86.55)
MAPE Improvement	-88.1% (12.27 pp)	-88.9% (13.35 pp)	-84.0% (8.71 pp)
Classification Accuracy	+64.86 pp	+64.59 pp	+51.26 pp
Directional Accuracy	+83.34 pp	+91.67 pp	+91.67 pp
Inference Latency	+27% (overhead)	+21% (overhead)	+21% (overhead)

Note: pp = percentage points. Inference latency is the only dimension on which MACPGANA does not lead; all other dimensions show substantial MACPGANA superiority.

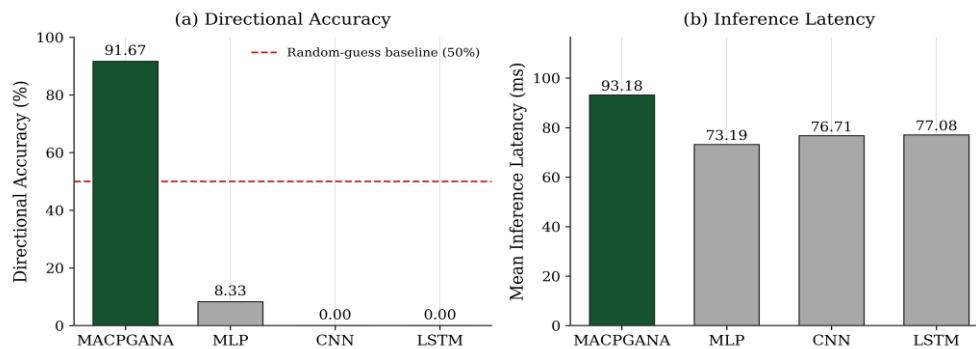


Fig. 4. Comparison of (a) directional accuracy and (b) mean inference latency across MACPGANA and the three baseline models.

The negative R^2 values for MACPGANA (-0.58) and MLP (-46.31) reflect the narrow approximately ₹200 price range of the test partition relative to training-partition variance, a well-documented statistical artefact in small-sample evaluation of low-volatility test windows, rather than a failure of either model; absolute error metrics (RMSE, MAE, MAPE) remain the appropriate primary performance indicators in this setting, and on all three, MACPGANA's advantage over every baseline is substantial and consistent.

7. Discussion

7.1 Why MACPGANA Outperforms the Baselines

Three architectural mechanisms plausibly account for MACPGANA's performance advantage, each addressing a specific limitation of the baseline architectures evaluated here. First, comprehensive multimodal feature integration: the baseline models access only raw price sequences or basic engineered features, a single perspective on price dynamics, whereas MACPGANA's feature pipeline provides simultaneous, complementary views of market momentum, periodic structure, and autoregressive pattern, which plausibly explains the 91.67% directional accuracy that sequence-only architectures fail to achieve. Second, GAN-based adaptive feature selection reduces the super feature vector to a smaller, high-predictive-value subset by evaluating candidate subsets through adversarial training rather than individual correlation thresholds, a globally informed selection strategy that limits the spurious correlations to which small training sets are prone. Third, the autoencoder's 16-dimensional bottleneck compresses the selected features in a manner that constrains the model toward generalisable rather than memorised representations, while the Q-Learning component provides continuous post-training adaptation without requiring full retraining.

7.2 Relation to Existing Literature

The relative ranking observed among the three baselines — LSTM outperforming CNN, which in turn outperforms MLP is consistent with prior Indian agricultural price-forecasting studies reporting LSTM's advantage over feedforward and regression baselines for crop price series [8,24]. The CNN-over-MLP improvement observed here is similarly consistent with prior work documenting CNN's advantage for extracting short-horizon local temporal patterns in financial and agricultural time series [10]. The benefit attributable to GAN-based feature/representation selection is consistent with Diqi, Hiswati, and Nur's [39] finding that GAN-based architectures outperform traditional methods under volatile price conditions, despite their non-agricultural setting. Taken together, the present results extend this literature by showing that these individually reported advantages compound when integrated into a single hybrid framework, and critically that the resulting advantage is not an artefact of large-scale, multimodal training data, since it persists under the constrained single-commodity conditions evaluated here.

7.3 Practical Implications

For farmers and producer organisations, the combination of 94.59% classification accuracy and 91.67% directional accuracy support evidence-based decisions on harvest timing, storage allocation, and market-entry timing, with a substantially tighter error margin ($\pm ₹59$, approximately) than any baseline ($\pm ₹215-327$). For traders and market intermediaries, directional accuracy of 91.67% supports confident inventory adjustment ahead of anticipated price movements, with the 93.18 ms inference latency permitting same-day strategy adjustment as new information arrives. For policymakers, forecast reliability at this level is sufficient to inform the timing of procurement or strategic-reserve interventions. The 21–27% latency overhead relative to the baselines is operationally immaterial given that agricultural commodity prices update at most once daily, well above the 93.18 ms prediction time required.

7.4 Limitations

Five boundary conditions qualify the generalisability of these findings. First, the evaluation covers a single commodity, market, and calendar year; validation across multiple commodities, regions, and years is needed to establish broader applicability. Second, the 36-sample test set introduces non-trivial statistical uncertainty around point estimates. Third, the current implementation relies solely on historical price data and price-derived engineered features, without exogenous variables such as meteorological data, transportation costs, or policy signals known to influence agricultural price formation. Fourth, only next-day ($t+1$) prediction is evaluated; performance over the longer 7–30-day horizons relevant to planting and storage decisions remains uncharacterised. Fifth, the reported inference latency assumes centralised server deployment; edge deployment on resource-constrained devices may require model compression.

8. Conclusion

This study evaluated MACPGANA, a hybrid GAN–autoencoder–Q-Learning framework for agricultural commodity price forecasting, against MLP, CNN, and LSTM baselines under a controlled, single-commodity, data-constrained setting using 2018 Raipur tomato price data. MACPGANA achieved a MAPE of 1.66% and a normalised RMSE of 0.0128 (₹29.38/quintal), corresponding to MSE reductions of 93.0–97.0% relative to the baselines; under a $\pm 10\%$ tolerance-band criterion, it achieved 94.59% classification accuracy and an F1-score of 97.22%, against 29.73–43.33% accuracy for the baselines; and it achieved 91.67% directional accuracy against 0–8.33% for the baselines, at a modest 21–27% inference-latency overhead. These results demonstrate that MACPGANA’s architectural advantages multimodal feature integration, adversarial feature selection, and adaptive Q-Learning optimisation are not contingent on the large-scale, multimodal training conditions of its original validation, and that the framework retains substantial predictive advantage when applied to the single-commodity, short-history conditions typical of regional agricultural markets in practice.

9. Future Work

Future work should extend this single-commodity validation along five directions identified by the limitations in Section 7.4: (i) replication across multiple commodities, markets, and growing seasons to establish the generality of the present findings; (ii) incorporation of exogenous predictors : meteorological data, transportation cost indices, and policy announcements into the multimodal feature pipeline; (iii) extension to multi-step forecasting horizons (7–30 days) relevant to planting and storage decisions; (iv) formal uncertainty quantification, such as bootstrap confidence intervals or conformal prediction, to characterise estimate reliability under small test-set conditions; and (v) evaluation of model compression techniques (e.g., knowledge distillation, parameter pruning) to support edge deployment on resource-constrained devices used in regional agricultural advisory systems.

Acknowledgement

The author expresses special thanks to the Government of India and the Horticulture Department for providing the necessary data on agriculture commodity prices. Special appreciation is also extended to Dr. Ramkrishna Rathore for guidance and valuable suggestions throughout the development of the research paper.

Conflict of Interest

The author declares no conflict of interest.

Data Availability

The dataset analysed in this study is publicly available from the National Horticulture Board, Government of India (<https://www.nhb.gov.in/OnlineClient/MISDailyReport.aspx>).

References

1. Ministry of Agriculture and Farmers Welfare, Annual Report 2022–23, Government of India, New Delhi, 2023.
2. H. C. J. Godfray et al., “Food security: the challenge of feeding 9 billion people,” *Science*, vol. 327, no. 5967, pp. 812–818, 2010.
3. C. L. Gilbert and C. W. Morgan, “Food price volatility,” *Philosophical Transactions of the Royal Society B*, vol. 365, no. 1554, pp. 3023–3034, 2010.
4. G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Hoboken, NJ, USA: Wiley, 2015.
5. Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
6. S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
7. R. Kambo, J. Pillai, S. Soni, and G. Urkude, “MACPGANA: design of a highly efficient multimodal agriculture commodity price prediction model via generative adversarial networks & autoencoders,” *International Journal of Information Technology*, 2025, doi: 10.1007/s41870-025-02425-z.
8. R. K. Paul et al., “Machine learning techniques for forecasting agricultural prices: a case of brinjal in Odisha, India,” *PLoS ONE*, 2022, doi: 10.1371/journal.pone.0270553.
9. S. Banerjee and A. C. Mondal, “An ingenious method for estimating future crop prices that emphasises machine learning and deep learning models,” *International Journal of Information Technology*, vol. 15, no. 8, pp. 4291–4313, 2023.
10. S. Verma, S. P. Sahu, and T. P. Sahu, “Discrete wavelet transform-based feature engineering for stock market prediction,” *International Journal of Information Technology*, vol. 15, no. 2, pp. 1179–1188, 2023.
11. R. Dharavath and E. Khosla, “Seasonal ARIMA to forecast fruits and vegetable agricultural prices,” in *Proc. IEEE Int. Symp. Smart Electronic Systems (iSES)*, 2019, pp. 47–52.

12. N. Bakhtadze, E. Maximov, and N. Maximova, "Local wheat price prediction models," in Proc. IEEE 7th Int. Conf. Control Science and Systems Engineering (ICCSSE), 2021, pp. 223–227.
13. G. Hegde, V. R. Hulipalled, and J. B. Simha, "A study on agriculture commodities price prediction and forecasting," in Proc. Int. Conf. Smart Technologies in Computing, Electrical and Electronics (ICSTCEE), 2020.
14. F. Sun, X. Meng, Y. Zhang, Y. Wang, H. Jiang, and P. Liu, "Agricultural product price forecasting methods: a review," *Agriculture*, vol. 13, no. 9, 2023.
15. A. Dutta, A. Nayak, R. Aditya, R. Panda, and N. K. Nagwani, "A neuro fuzzy system based inflation prediction of agricultural commodities," in Proc. Int. Conf. Computing, Communication and Networking Technologies (ICCCNT), 2020.
16. D. Zhang, S. Chen, L. Liwen, and Q. Xia, "Forecasting agricultural commodity prices using model selection framework with time series features and forecast horizons," *IEEE Access*, vol. 8, pp. 28197–28209, 2020.
17. Y. Liu and Y. Lv, "Commodity price evaluation based on improved data mining methods," in Proc. Int. Conf. E-Commerce and Internet Technology (ECIT), 2020, pp. 145–148.
18. M. M. Rashid, B. S. Bari, Y. Yusup, M. A. Kamaruddin, and N. Khan, "A comprehensive review of crop yield prediction using machine learning approaches with special emphasis on palm oil yield prediction," *IEEE Access*, vol. 9, pp. 63406–63439, 2021.
19. G. Hegde, V. R. Hulipalled, and J. B. Simha, "Price prediction of agriculture commodities using machine learning and NLP," in Proc. 2nd Int. Conf. Smart Technologies in Computing, Electrical and Electronics (ICSTCEE), 2021.
20. S. K. Sharma, D. P. Sharma, and J. K. Verma, "Study on machine-learning algorithms in crop yield predictions specific to Indian agricultural contexts," in Proc. Int. Conf. Computational Performance Evaluation (ComPE), 2021, pp. 155–166.
21. B. Adhikari et al., "Building market intelligence systems for agricultural commodities: a case study based on cardamom," in Proc. 2nd Int. Conf. Secure Cyber Computing and Communications (ICSCCC), 2021, pp. 12–16.
22. Haviluddin et al., "A backpropagation neural network algorithm in agricultural product prices prediction," in Proc. 3rd East Indonesia Conf. Computer and Information Technology (EIConCIT), 2021, pp. 399–404.
23. G. Urkude and M. Pandey, "AgriSense: automatic irrigation utility system using wireless sensor network and web of things," in Proc. 2nd Int. Conf. Advanced Computational and Communication Paradigms (ICACCP), 2019.
24. T. Banerjee, S. Sinha, and P. Choudhury, "Long term and short term forecasting of horticultural produce based on the LSTM network model," *Applied Intelligence*, vol. 52, no. 8, pp. 9117–9147, 2022.
25. A. M. Joshi and S. Patel, "A CNN-Bidirectional LSTM approach for price forecasting of agriculture commodities in Gujarat," in Proc. Int. Conf. Applied Artificial Intelligence and Computing (ICAAIC), 2022, pp. 266–272.
26. R. L. Manogna and A. K. Mishra, "Forecasting spot prices of agricultural commodities in India: application of deep-learning models," *Intelligent Systems in Accounting, Finance and Management*, 2021.
27. Y. Zhang and S. Na, "A novel agricultural commodity price forecasting model based on fuzzy information granulation and MEA-SVM model," *Mathematical Problems in Engineering*, vol. 2018, 2018.
28. S. Gupta et al., "Machine learning- and feature selection-enabled framework for accurate crop yield prediction," *Journal of Food Quality*, vol. 2022, 2022.
29. A. Vohra, N. Pandey, and S. K. Khatri, "Decision making support system for prediction of prices in agricultural commodity," in Proc. Amity Int. Conf. Artificial Intelligence (AICAI), 2019, pp. 345–348.
30. A. K. Mahto, M. A. Alam, R. Biswas, J. Ahmed, and S. I. Alam, "Short-term forecasting of agriculture commodities in context of Indian market for sustainable agriculture by using the artificial neural network," *Journal of Food Quality*, vol. 2021, 2021.
31. Y. Fang, B. Guan, S. Wu, and S. Heravi, "Optimal forecast combination based on ensemble empirical mode decomposition for agricultural commodity futures prices," *Journal of Forecasting*, vol. 39, no. 6, pp. 877–886, 2020.
32. K. Kurumatani, "Time series forecasting of agricultural product prices based on recurrent neural networks and its evaluation method," *SN Applied Sciences*, vol. 2, no. 8, 2020.
33. M. R. C. Van Oordt, P. A. Stork, and C. G. De Vries, "On agricultural commodities' extreme price risk," *Extremes*, vol. 24, no. 3, pp. 531–563, 2021.
34. H. Ouyang, X. Wei, and Q. Wu, "Agricultural commodity futures prices prediction via long- and short-term time series network," *Journal of Applied Economics*, vol. 22, no. 1, pp. 468–483, 2019.
35. H. A. Marfatia, Q. Ji, and J. Luo, "Forecasting the volatility of agricultural commodity futures: the role of co-volatility and oil volatility," *Journal of Forecasting*, vol. 41, no. 2, pp. 383–404, 2022.
36. J. Crespo Cuaresma, J. Hlouskova, and M. Obersteiner, "Agricultural commodity price dynamics and their determinants: a comprehensive econometric approach," *Journal of Forecasting*, 2021.
37. X. Xu and Y. Zhang, "Commodity price forecasting via neural networks for coffee, corn, cotton, oats, soybeans, soybean oil, sugar, and wheat," *Intelligent Systems in Accounting, Finance and Management*, vol. 29, no. 3, pp. 169–181, 2022.
38. X. He, M. Carriquiry, A. Elobeid, D. Hayes, M. Li, and W. Zhang, "China's corn and biofuel policies and agricultural trade: projections from an international agricultural commodity market model," *Agribusiness*, vol. 38, no. 4, pp. 970–989, 2022.
39. M. Diqi, M. E. Hiswati, and A. S. Nur, "StockGAN: robust stock price prediction using GAN algorithm," *International Journal of Information Technology*, vol. 14, no. 5, pp. 2309–2315, 2022.

40. L. Goel and A. Mishra, "A survey of recent deep learning algorithms used in smart farming," in Proc. IEEE Region 10 Symposium (TENSYMP), 2022, pp. 1–6.
41. Z. Hajirahimi, M. Khashei, and A. Z. Hamadani, "Principal component-based hybrid model for time series forecasting," International Journal of Information Technology, vol. 15, no. 6, pp. 3045–3053, 2023.
42. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
43. I. Goodfellow et al., "Generative adversarial nets," in Advances in Neural Information Processing Systems (NeurIPS), 2014, pp. 2672–2680.