International Journal of Enhanced Research in Science, Technology & Engineering
ISSN: 2319-7463, Vol. 15 Issue 10, October-2026

MandiSense AI: Real-Time Price Intelligence and Predictive Market Optimization System

Monika R[1], Kanchana K R[2], Jonah Davis[3], Shivraj Soudi[4], Prof. Amith Goutham[5]

[1,2,3,4] Department of Computer Science and Engineering with Data Science, AMC Engineering College, Bannerghatta, Karnataka, India
[5](Prof. of CSE-DS) Department of Computer Science and Engineering with Data Science, AMC Engineering College, Bannerghatta, Karnataka, India

Corresponding author, Monika R, e-mail: monikamonim526@gmail.com

ABSTRACT

Farmers in India sell many crops through regulated wholesale markets called mandis. Prices change daily and differ between mandis, but farmers often lack timely and organised price information, so they may sell at a lower price or choose a mandi without considering transport cost. This paper presents MandiSense AI, a web-based platform that combines up-to-date (latest available) mandi price information, multi-mandi comparison, future price prediction, best-mandi recommendation, transport cost and farmer profit calculation, price alerts, and anomaly detection for suspicious price entries. The system loads mandi price data from stored CSV files (a synthetic sample dataset designed using the structure and field definitions of publicly available government mandi-price information), cleans and stores it, and applies simple, explainable methods (moving average and linear trend, compared with naive and seasonal baselines) to forecast short-term prices. Mandis are ranked by estimated net profit, defined as sale revenue minus transport and other costs. The system was tested on a synthetic sample dataset of 10,080 monthly records covering 14 commodities and 15 mandis in Karnataka (January 2023 to December 2025). In a rolling one-month-ahead test on 2025 data, a seasonal-naive baseline achieved a MAPE of 3.69% and a naive baseline 5.19%, while a 3-month moving average and a linear trend achieved 7.80% and 8.72%. The paper also discusses limitations and future extensions.

Keywords: **agricultural markets, mandi prices, price forecasting, time series, recommendation system, anomaly detection, transport cost, farmer decision support, data science**

INTRODUCTION

Background

Agriculture is a major source of livelihood in India, and a large share of farm produce is sold through mandis, which are wholesale markets where buyers and sellers meet. The price a farmer receives depends on the crop, its quality, the day of sale, the quantity arriving in the market and the location of the mandi. Prices for the same crop can differ noticeably between nearby markets on the same day.

Government platforms such as Agmarknet (Ministry of Agriculture and Farmers Welfare, Government of India, n.d.) and the National Agriculture Market (eNAM; Government of India, n.d.-a) publish market information and support electronic trading. However, raw price listings are often difficult for an individual farmer to use directly. A farmer needs practical answers: which mandi gives the best return after transport, whether the price will rise or fall, and whether a price entry is reliable.

Studies in development economics show that access to market information can affect market outcomes. Work on fisheries in Kerala (Jensen, 2007), grain markets in Niger (Aker, 2010) and farmers in central India (Goyal, 2010) examined how information access relates to price dispersion, waste and welfare. These studies motivate accessible price information systems, although they do not describe the system proposed here.

Problem Statement

Farmers often decide where and when to sell using limited or delayed information, such as the price quoted by a local trader. Existing portals show prices but generally do not (a) compare mandis by net profit after transport, (b) give future price estimates, (c) notify the farmer when a target price is reached, or (d) warn about unusual or possibly incorrect

entries. The problem addressed here is how to build an integrated, easy-to-use system that gives reliable current prices, short-term forecasts and a clear recommendation of where to sell, while checking data quality.

Objectives

1. To collect and present up-to-date mandi prices for selected crops and markets.
2. To compare prices across multiple mandis.
3. To predict short-term future prices using machine learning or time-series models.
4. To recommend the best mandi using price, transport cost and other costs.
5. To calculate transport cost and expected farmer profit.
6. To send alerts when a farmer's price condition is met.
7. To detect anomalous or possibly manipulated price entries.

Organisation of the Paper

Section 2 reviews related work and existing systems. Section 3 presents the proposed system and Section 4 its architecture. Section 5 describes the dataset and preprocessing. Section 6 explains the machine learning methodology and Section 7 the recommendation, transport cost and profit module. Section 8 describes the user interface and alerts. Section 9 discusses results, Section 10 lists advantages and limitations, Section 11 outlines future scope, and Section 12 concludes the paper.

RELATED WORK

Literature Review

Jensen (2007) studied mobile phone adoption among fishermen in Kerala and reported that access to information was linked to reduced price dispersion and less waste. Aker (2010) examined mobile phones and grain markets in Niger, and Goyal (2010) studied an internet-based information and procurement initiative for soybean farmers in central India. Together, these papers suggest that better market information can improve market functioning; readers should consult the originals for exact findings.

For forecasting, the ARIMA family is a classical approach described by Box et al. (2015) and Hyndman and Athanasopoulos (2021). Hybrid ARIMA and neural network models were proposed by Zhang (2003), and the Prophet model for trend and seasonality by Taylor and Letham (2018). Random Forest (Breiman, 2001), XGBoost (Chen & Guestrin, 2016) and LSTM networks (Hochreiter & Schmidhuber, 1997) are widely used learning methods, and scikit-learn (Pedregosa et al., 2011) provides implementations of many standard tools.

Chandola et al. (2009) surveyed anomaly detection techniques, and Isolation Forest (Liu et al., 2008) isolates unusual points using random trees. Hyndman and Koehler (2006) discuss forecast accuracy measures, which guide the metrics used in this work.

Several studies have forecast Indian agricultural prices. Manogna et al. (2025a) compared a GRU deep-learning model with ARIMA for onion and tomato prices and found much lower errors for the GRU model, and the same group combined LSTM and GARCH to forecast price volatility (Manogna et al., 2025b). Suresha et al. (2021) compared ARIMA, RNN, LSTM, VAR, Random Forest and XGBoost for retail onion, potato and tomato prices, using arrival volumes and diesel prices, and also looked at anomaly detection. Mitra and Paul (2021) applied an ARFIMA long-memory model to daily wholesale rice prices. Beyond forecast accuracy, Ma et al. (2018) built a web application that pulls Agmarknet data, fills missing values and shows interpretable price-direction forecasts for small farmers, and Ladhar et al. (2023) developed a price forecasting and sales recommendation system, with an Android app, for soyabean farmer cooperatives that suggests when to sell.

Research gap. Most of these studies focus on forecast accuracy or on when to sell. In the studies reviewed here we did not find one that combines multi-mandi comparison, transport-adjusted net profit, threshold alerts and anomaly flags in one simple, explainable tool, which is the aim of MandiSense AI. This gap comes from a limited review and should be checked against newer work.

Existing Systems

Farmers and traders currently use government portals such as Agmarknet (Ministry of Agriculture and Farmers Welfare, Government of India, n.d.) and eNAM (Government of India, n.d.-a), local traders, and general-purpose mobile applications. Ma et al. (2018) describe Agmarknet as providing price and arrival data for over a thousand markets through a website and mobile app, and list other apps such as FarmBee (prices, weather and news) and Kisan Suvidha (advisories); they knew of no app that gave interpretable price forecasts. These descriptions come from a 2018 paper, so the current features of these tools should be checked. Limitations that may apply to such tools include:

- Prices are mostly shown as lists with limited comparison.
- Future price estimates may not be provided.
- Transport cost and net profit are often not included.

- Personalised alerts and data-quality checks may be limited.

Proposed System

MandiSense AI is a web-based platform developed by the project team as an academic project, with the following modules:

1. **Real-time price module:** shows the latest available minimum, maximum and modal prices for a crop and mandi from the uploaded CSV dataset. Prices are as current as the most recent file loaded, so this module is not a live market feed. Data are refreshed only when a new CSV file is uploaded.
2. **Mandi comparison module:** shows prices of one crop across several mandis.
3. **Price prediction module:** forecasts prices for the next 7, 15 or 30 days.
4. **Recommendation module:** ranks mandis by expected net profit.
5. **Transport and profit module:** estimates transport cost and farmer profit.
6. **Price alert module:** notifies the user by in-app notification when a condition is met.
7. **Anomaly detection module:** flags unusual price records before use.

System Architecture

The system follows a modular layered architecture, shown in Figure 1 and summarised in Table 1. Raw prices enter through the ingestion layer, are cleaned and screened for anomalies, and are stored. The prediction and recommendation modules read the stored data, and the backend serves results to the web dashboard and the alert service.

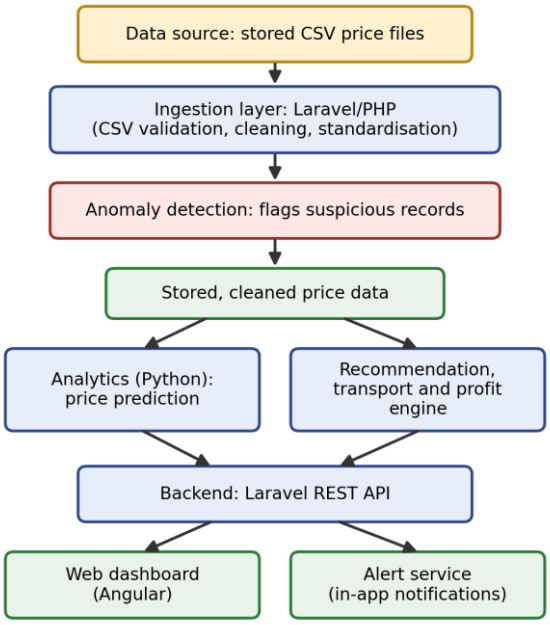


Figure 1: ***System Architecture of MandiSense AI***

Table 1: ***Architecture Layers***

Layer	Purpose	Technology
Data source	Raw mandi prices	Stored CSV files
Ingestion	Collect, clean, standardise	Laravel/PHP with CSV validation
Analytics	Prediction, recommendation, anomalies	Python/statistical methods
Backend	Serves results to the UI	Laravel REST API
Frontend	Dashboard and comparison	Angular
Alerts	Sends notifications	In-app notifications

DATASET AND PREPROCESSING

Table 2 summarises the dataset. The synthetic dataset contains 10,080 records covering 14 commodities and 15 mandis across Karnataka, with monthly observations from January 2023 to December 2025. For system development and testing, a synthetic sample dataset was used to provide a consistent and controlled set of mandi-price records. The dataset was designed using the structure and field definitions of publicly available government mandi-price information such as Agmarknet (Ministry of Agriculture and Farmers Welfare, Government of India, n.d.) and the Open Government

Data (OGD) Platform (Government of India, n.d.-b). The values in the dataset are intended for system testing and demonstration and should not be interpreted as official live market prices, so the results show how the system behaves and are not evidence of accuracy on real markets. The application accepts a CSV file with the required columns Date, State, District, Mandi, Commodity, Min_Price, Max_Price and Modal_Price, and the optional columns Variety and Arrivals_Tonnes. Prices are in Rs per quintal and arrivals are in tonnes. Any assumed values, such as vehicle rates, are stated as assumptions.

Table 2: ***Dataset Description***

Item	Details
Source	Synthetic sample data (not official prices), stored as a CSV file, in the Agmarknet/data.gov.in column format
Commodities	14: Brinjal, Cabbage, Cauliflower, Cotton, Green Chilli, Groundnut, Maize, Onion, Paddy, Potato, Ragi, Tomato, Turmeric, Wheat
Mandis / state	Karnataka; 15 APMC mandis across 20 districts (Bagalkot, Belagavi, Bengaluru, Bidar, Chamarajanagar, Chikkaballapur, Chikkamagaluru, Chitradurga, Dharwad, Hassan, Kolar, Mandya, Mysuru, Shivamogga, Tumakuru)
Date range	1 Jan 2023 to 1 Dec 2025 (36 monthly observations)
No. of records	10,080
Attributes	Date (month of observation), State, District, Mandi (APMC name), Commodity, Variety (optional), Min_Price, Max_Price and Modal_Price (Rs per quintal), Arrivals_Tonnes (optional; tonnes)
Licence / terms	Not applicable (synthetic data, not an official or third-party dataset)
Train / test split	Train: Jan 2023 to Dec 2024 (24 months); test: Jan to Dec 2025 (12 months)

The dataset was checked before use. No missing values and no exact duplicate rows were found, and every record satisfies Min_Price ≤ Modal_Price ≤ Max_Price. The Variety field contains only the value “Other”. The file lists 20 districts but only 15 distinct mandi names, so some mandi names appear under more than one district; records with the same mandi, commodity and date (2,520 cases) were averaged, giving 210 monthly price series (14 commodities × 15 mandis). Prices were not scaled, and no encoding was needed because each mandi-commodity pair is forecast as its own series.

MACHINE LEARNING METHODOLOGY

The methodology pipeline is shown in Figure 2.

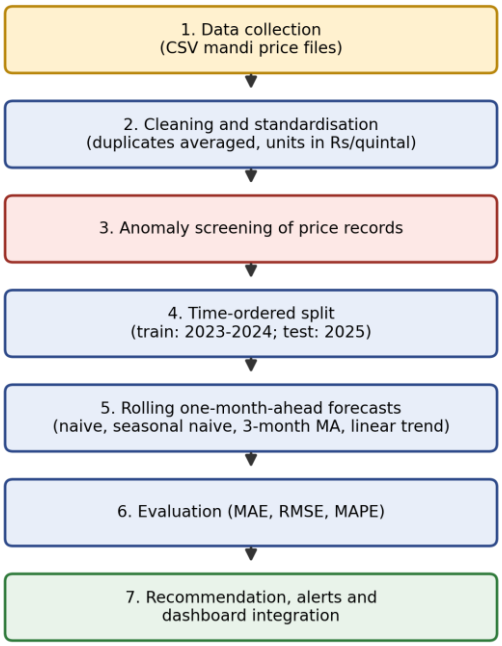


Figure 2: Methodology Pipeline of the Proposed System

Forecasting Models

Four simple, explainable methods were evaluated: (i) a naive baseline that repeats the previous month's modal price, (ii) a 3-month moving average, (iii) a linear trend fitted to all earlier observations, and (iv) a seasonal-naive baseline that repeats the price of the same month in the previous year. These are standard forecasting methods described by Hyndman and Athanasopoulos (2021). The Price Prediction module uses a 3-month moving average and a linear trend method for short-term price forecasting.

Features and Validation

Each series is forecast one month ahead in a rolling way, using only observations earlier than the month being predicted (an expanding window), so the test set never leaks into training. Only past modal prices were used; no other features such as arrivals or calendar variables were added. No cross-validation or hyperparameter tuning was performed, and the 3-month window was fixed in advance. Errors are reported as MAE, RMSE and MAPE over all 2,520 test forecasts (210 series × 12 months of 2025).

Evaluation Metrics

Forecast quality is measured by Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE; Hyndman & Koehler, 2006):

$$MAE = (1/n)\ \Sigma\ |y_i - \hat{y}_i| \quad (1)$$

$$RMSE = \sqrt{[(1/n)\ \Sigma\ (y_i - \hat{y}_i)^2]} \quad (2)$$

$$MAPE = (100/n)\ \Sigma\ |(y_i - \hat{y}_i)/y_i| \quad (3)$$

Anomaly Detection

Price records are screened using a fixed-percentage threshold on the percentage change between consecutive prices, together with consistency checks such as Min_Price ≤ Modal_Price ≤ Max_Price. Flagged records are marked for review instead of being silently deleted. In this paper, “manipulated data” means entries that are statistically unusual or inconsistent; the system cannot prove intentional manipulation. In the evaluation in Section 9, the rule was a z-score on the monthly percentage change with a threshold of $|z| > 3$ within each series, tested by injecting artificial spikes.

Recommendation, Transport Cost and Profit

For each candidate mandi *m*, the net profit is calculated as

$$NetProfit(m) = (P_m \times Q) - T_m - O \quad (4)$$

where P_m is the expected price per unit (current or predicted), Q is the quantity, T_m is the transport cost and O is other cost such as loading or market fees. In the best-mandi recommendation, only transport cost is deducted (O = 0); the transport and profit calculator also lets the user enter other expenses. Quantity is converted from tonnes to quintals (1 tonne = 10 quintals) so it matches prices quoted per quintal. Transport cost is estimated as

$$T_m = D_m \times R \quad (5)$$

where D_m is the distance to the mandi and R is the transport cost per km entered by the user. Distance is calculated between the farmer's location and the selected mandi. Mandis are then ranked in descending order of net profit, following the workflow in Supplementary Figure S1 (Supplementary Data file).

User Interface and Price Alerts

The dashboard lets the user filter prices by state, district, mandi, commodity and date. It shows summary cards for the latest date, lowest, highest and average modal price, total arrivals and number of commodities, along with a commodity price chart and a mandi-wise price chart (Figure 3). A dataset page allows the user to upload a CSV file, download a sample CSV and preview the loaded records.

The recommendation screen shows transport cost, profit and the recommended mandi (Supplementary Figure S2). Users set an alert when the price of a crop at a chosen mandi reaches a target value, and flagged anomalous records are shown for review (Supplementary Figure S3). Alert conditions are checked against the latest dataset price when the user runs an alert check, and a triggered alert is shown as an in-app notification (no SMS or e-mail).

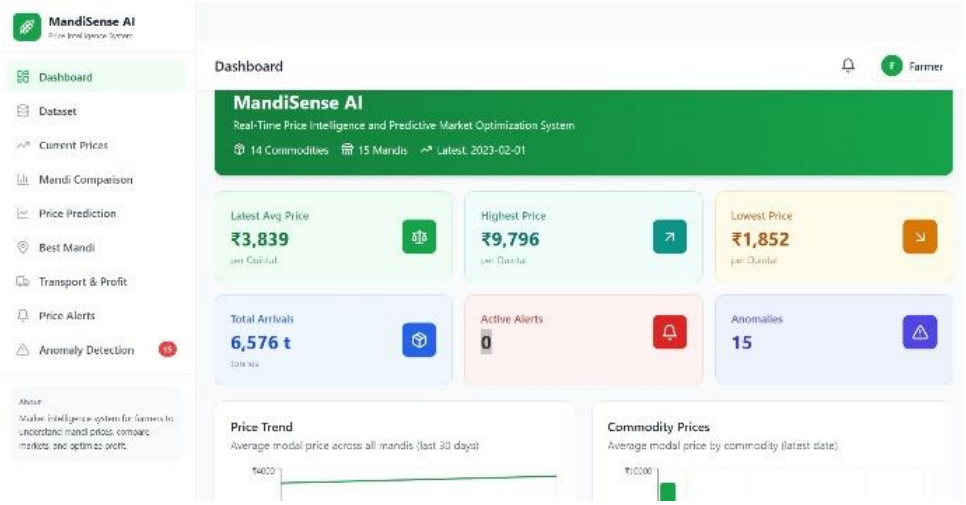


Figure 3: ***Dashboard With Latest Prices and Mandi Comparison***

Note. The screenshot shows the dashboard with a latest date of 1 February 2023, so its values differ from the December 2025 summary in Table 4.

RESULTS AND DISCUSSION

Forecasting Performance

Table 3 averages the one-month-ahead errors over all 210 mandi-commodity series for the 12 months of 2025 (2,520 forecasts per method; modal price in Rs per quintal). The seasonal-naive baseline was the most accurate (MAPE 3.69%), followed by the naive baseline (5.19%). The 3-month moving average (7.80%) and the linear trend (8.72%) were less accurate than simply repeating last month's price, probably because they lag behind or over-extrapolate a series that follows a strong yearly cycle. These figures come from a synthetic dataset with a repeating seasonal pattern and should not be read as accuracy on real market prices. Figure 4 shows an example series.

Table 3: ***Forecast Error on the Test Set***

Crop	Mandi	Model	MAE	RMSE	MAPE (%)
All (14)	All (15)	Naive baseline	180.52	250.11	5.19
All (14)	All (15)	Seasonal naive (same month last year)	128.71	187.31	3.69
All (14)	All (15)	Moving average (3 months)	271.34	353.73	7.80
All (14)	All (15)	Linear trend (all prior data)	305.04	397.86	8.72

Note. MAE and RMSE are in Rs per quintal. MAPE = mean absolute percentage error.

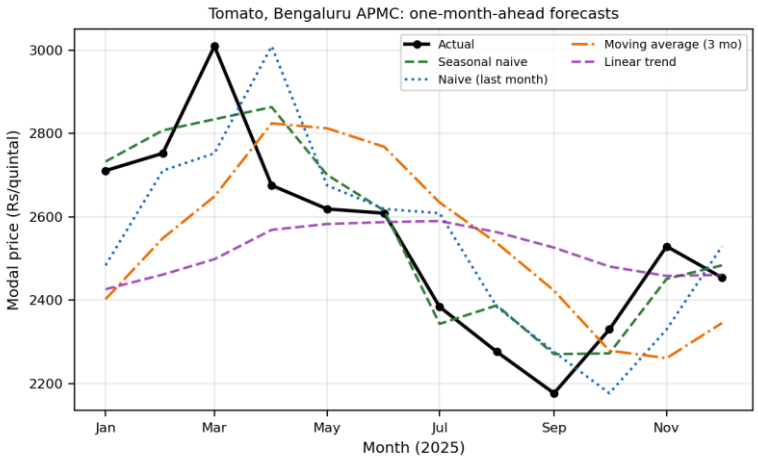


Figure 4: Actual Versus Predicted Prices on the Test Set (Tomato, Bengaluru APMC, 2025)

Recommendation and Anomaly Detection

Worked example (illustrative; distances and the transport rate are assumed, not measured). A farmer near Kolar has 5 tonnes (50 quintals) of tomato and compares five mandis using the December 2025 modal prices in the synthetic dataset, with an assumed rate of Rs 15 per km. Kolar APMC (5 km, Rs 2,405.16 per quintal) gives a net profit of Rs 120,183; Chikkaballapur APMC (55 km, Rs 2,448.18) gives Rs 121,584; Bengaluru APMC (70 km, Rs 2,454.00) gives Rs 121,650; Tumakuru APMC (150 km, Rs 2,566.08) gives Rs 126,054; and Chitradurga APMC (240 km, Rs 2,920.38) gives Rs 142,419 (revenue Rs 146,019 minus transport Rs 3,600). Chitradurga APMC is recommended. Here it is also the highest-price mandi, because its price advantage outweighs its extra transport cost.

Anomaly check. A z-score rule on the monthly percentage change ($|z| > 3$ within each mandi-commodity series) flagged 5 of 7,350 price changes in the original data, all increases. To test it, 100 artificial spikes of 40% (up or down) were injected into the series; the rule flagged 40 of them (recall 0.40) and made 2 other false flags (precision 0.95). With spikes of 30% and 50%, recall was 0.32 and 0.31. Recall is low because the synthetic data already have large seasonal swings, so moderate spikes stay inside the normal range; a seasonal or rolling baseline would likely detect more (Figure 5).

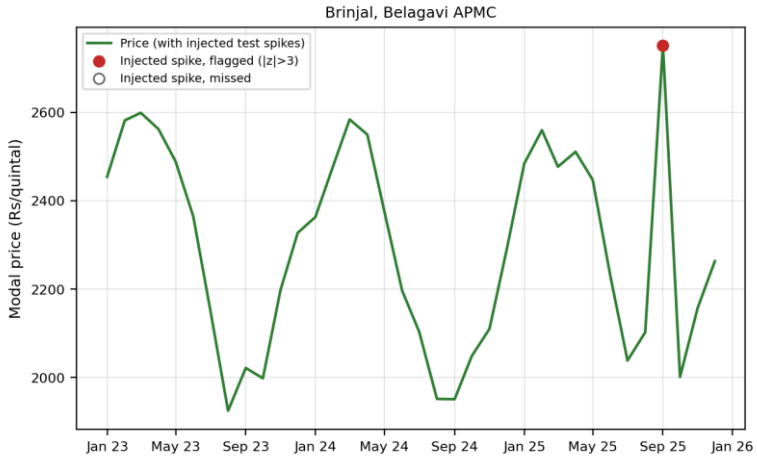


Figure 5: Anomalies Flagged on a Price Series (Brinjal, Belagavi APMC, With Injected Test Spikes)

Module Testing

Table 4 summarises the module tests.

Table 4: ***Summary of Module Testing***

Module	Test performed	Outcome
Real-time prices	Latest-month summary computed from the CSV (script check)	Latest date 1 Dec 2025; average modal price Rs 3,492.28/quintal; highest Rs 9,406.56; lowest Rs 1,652.02; total arrivals 8,300.5 tonnes (280 records)
Mandi comparison	Tomato modal prices across 15 mandis, latest month (script check)	Lowest: Belagavi APMC Rs 2,269.74; highest: Chitradurga APMC Rs 2,920.38
Prediction	One-month-ahead test on 2025 data (Table 3)	MAPE from 3.69% (seasonal naive) to 8.72% (linear trend)
Recommendation	Worked example in Section 9.2 (assumed distances and rate)	Chitradurga APMC recommended; net profit Rs 142,419
Transport / profit	Hand-checked: 5 tonnes = 50 quintals at Rs 2,920.38, 240 km at Rs 15/km	Gross Rs 146,019; transport Rs 3,600; profit Rs 142,419 (Rs 2,848.38 per quintal)
Alerts	Rule check: Tomato, Bengaluru APMC, alert "price goes above Rs 2,400"	Latest price Rs 2,454.00, so the alert condition is met
Anomaly detection	100 injected 40% spikes; rule $\|z\| > 3$ on monthly change	Recall 0.40, precision 0.95; 5 of 7,350 original changes flagged

In the worked example, transport cost did not change the recommendation: the highest-price mandi also gave the highest net profit, because its price was about Rs 515 per quintal higher than Kolar (about Rs 25,750 for 50 quintals) while its extra transport cost was about Rs 3,500. Transport cost would change the ranking only when price gaps are smaller than the extra transport cost. Sudden market events are not present in the data and were not tested. Formal user evaluation was not performed. The synthetic data follow an almost identical yearly price cycle (for example, the average Tomato price has a correlation of about 0.99 between consecutive years), which explains why the seasonal baseline performs best.

Advantages and Limitations

Advantages

- Combines prices, comparison, forecasts, recommendation, costs and alerts in one place.
- Recommends by net profit, which is more useful than the highest price alone.
- Uses simple, widely available and reproducible methods.
- Data-quality checks reduce the effect of unreliable entries.
- Alerts save the farmer from checking prices repeatedly.

Limitations

- Prediction quality depends on data completeness and accuracy.
- Forecasts cannot account for sudden events such as weather shocks or policy changes unless such data are added.
- Anomaly detection flags unusual values but cannot confirm manipulation.
- Transport cost is an estimate based on assumed rates. Rates and distances are entered or assumed values, not measured road costs.
- The system was tested on limited crops, mandis and dates. It used a synthetic sample dataset of 14 commodities and 15 mandis in Karnataka with monthly data from January 2023 to December 2025, so results may differ on real market data. Because the data are monthly, the 7-, 15- and 30-day forecast options in the app were not validated. The dataset also lists some mandi names under more than one district.
- No large-scale field trial with farmers was conducted.

Future Scope

- More crops, mandis and states.
- Weather, rainfall, arrival quantity and fuel price as extra features.
- Advanced models such as LSTM (Hochreiter & Schmidhuber, 1997) or hybrid models (Zhang, 2003), compared systematically.
- Regional language, voice and SMS support for low-connectivity users.
- Vehicle availability and dynamic transport pricing.
- Field testing with farmer groups and integration with official platforms such as eNAM (Government of India, n.d.-a), subject to their terms.

CONCLUSION

This paper presented MandiSense AI, a system that supports farmers' selling decisions through mandi price information, multi-mandi comparison, price prediction, net-profit-based recommendation, transport and profit calculation, price alerts and anomaly detection. In a one-month-ahead test on 2025 synthetic data, a seasonal-naive baseline gave a MAPE of 3.69%, while a 3-month moving average and a linear trend gave 7.80% and 8.72%. Although the current version has limits in data coverage, accuracy and field validation, it provides a base that can be extended with more data, better models and real-world testing.

ACKNOWLEDGMENT

The authors thank their guide, Prof. Amith Goutham, and the Department of Computer Science and Engineering (Data Science), AMC Engineering College, Bengaluru, for their guidance and support.

REFERENCES

1. Aker, J. C. (2010). Information from markets near and far: Mobile phones and agricultural markets in Niger. *American Economic Journal: Applied Economics, 2*(3), 46–59. https://doi.org/10.1257/app.2.3.46
2. Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control* (5th ed.). Wiley.
3. Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32. https://doi.org/10.1023/A:1010933404324
4. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882

5. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785
6. Government of India. (n.d.-a). *National Agriculture Market (eNAM)*. https://enam.gov.in
7. Government of India. (n.d.-b). *Open Government Data (OGD) Platform India*. https://data.gov.in
8. Goyal, A. (2010). Information, direct access to farmers, and rural market performance in central India. *American Economic Journal: Applied Economics, 2*(3), 22–45. https://doi.org/10.1257/app.2.3.22
9. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
10. Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts. https://otexts.com/fpp3/
11. Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting, 22*(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001
12. Jensen, R. (2007). The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics, 122*(3), 879–924. https://doi.org/10.1162/qjec.122.3.879
13. Ladhar, R., Sharma, S., Tangirala, S., Gupta, N., Azeem, A., Jain, A., Katakam, B. C., Bommakanti, A., Sankaraiah, C., Piridi, H. P., Yadav, K., Vittalrao, K., Rahman, M., Chutani, R., Shah, R., Katepallewar, R., Chakraborty, D., & Seth, A. (2023). AI-based market intelligence systems for farmer collectives: A case study from India. *ACM Journal on Computing and Sustainable Societies, 1*(1), 1–32. https://doi.org/10.1145/3609262
14. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *Proceedings of the 8th IEEE International Conference on Data Mining* (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17
15. Ma, W., Nowocin, K., Marathe, N., & Chen, G. H. (2018). *An interpretable produce price forecasting system for small and marginal farmers in India using collaborative filtering and adaptive nearest neighbors* (arXiv:1812.05173). arXiv. https://arxiv.org/abs/1812.05173
16. Manogna, R. L., Dharmaji, V., & Sarang, S. (2025a). Enhancing agricultural commodity price forecasting with deep learning. *Scientific Reports, 15*, Article 20903. https://doi.org/10.1038/s41598-025-05103-z
17. Manogna, R. L., Dharmaji, V., & Sarang, S. (2025b). A novel hybrid neural network-based volatility forecasting of agricultural commodity prices: Empirical evidence from India. *Journal of Big Data, 12*. https://doi.org/10.1186/s40537-025-01131-8
18. Ministry of Agriculture and Farmers Welfare, Government of India. (n.d.). *Agmarknet: Agricultural Marketing Information Network*. https://agmarknet.gov.in
19. Mitra, D., & Paul, R. K. (2021). Forecasting of price of rice in India using long-memory time-series model. *National Academy Science Letters, 44*, 289–293. https://doi.org/10.1007/s40009-020-01002-1
20. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830. https://jmlr.org/papers/v12/pedregosa11a.html
21. Suresha, H. P., Uppaladinni, M., & Tiwari, K. K. (2021). Indian commodity market price comparative study of forecasting methods: A case study on onion, potato and tomato. *Asian Journal of Research in Computer Science, 12*(4).
22. Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician, 72*(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080
23. Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing, 50*, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0