Agri Articles
(e-Magazine for Agricultural Articles)
Volume: 06, Issue: 04 (JULY-AUG, 2026)
Available online at http://www.agriarticles.com

When Artificial Intelligence Predicts the Harvest: Is Accuracy Enough?

*Dr. Chiranjit Mazumder[1], Nithyashree, M. L[2] and Utkarsh Tiwari[1]
[1]Scientist, Division of Agricultural Economics, ICAR-Indian Agricultural Research Institute, New Delhi, India
[2]SeniorScientist, Division of Agricultural Economics, ICAR-Indian Agricultural Research Institute, New Delhi, India
*Corresponding Author's email: majumder.chira@gmail.com

Artificial intelligence can combine satellite images, weather records, soil information and past yields to forecast a crop before it is harvested. The promise is enormous: earlier production estimates, better procurement planning, improved insurance assessment and more timely farm advice. Yet a model can score impressively on an accuracy test and still fail the farmer. India therefore needs to judge harvest forecasts by a broader standard: Do they work locally, reveal uncertainty, explain their reasoning, include diverse farmers and lead to better decisions?

A precise number can still be a poor forecast

Imagine that an artificial intelligence (AI) system predicts a wheat yield of 4.7 tonnes per hectare for a district. The number looks authoritative. It may appear on a dashboard with a map shaded green and a confidence score above 90%. But the district contains irrigated farms near a canal, rainfed plots on light soils, early- and late-sown wheat, and farmers using very different seed and fertilizer practices. Which of these farms does 4.7 tonnes describe?

This is the central problem in AI-based harvest prediction. Accuracy is essential, but it is not sufficient. A forecast is valuable only when it is accurate for the relevant place and time, honest about uncertainty, understandable to users and connected to a decision. For a statistician, a small prediction error is good. For a farmer, procurement agency or insurer, the more important question is whether the forecast improves an action without creating avoidable harm.

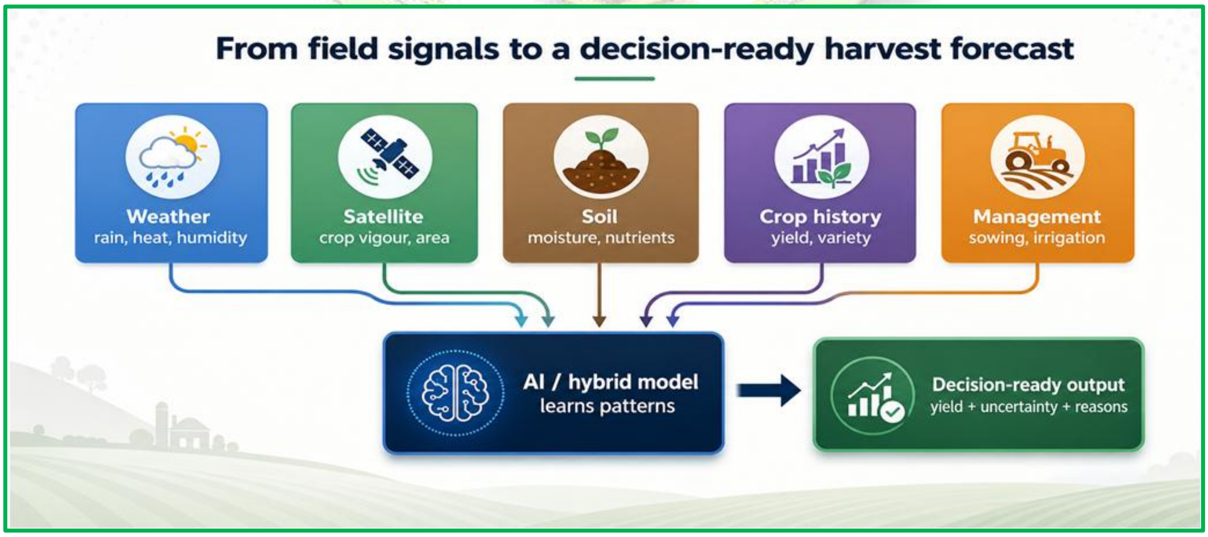

Figure 1. An original conceptual illustration of an AI-enabled harvest prediction pipeline.

How does AI see a harvest?

A crop leaves many digital traces before it reaches the threshing floor. Rainfall, temperature and humidity describe the season. Satellite images reveal crop area, greenness and stress. Soil information indicates moisture and nutrient conditions. Records of varieties, sowing dates, irrigation, pests and previous yields add local context. Machine-learning models search these data for patterns that are associated with final production. Reviews of crop-yield research show that temperature, rainfall and soil type are among the most frequently used inputs, while neural networks and other machine-learning methods are widely applied (van Klompenburg et al., 2020).

India is not beginning from zero. The Indian Space Research Organisation developed remote-sensing approaches for crop acreage and production estimation, later expanded through FASAL - Forecasting Agricultural Output using Space, Agro-meteorology and Land-based Observations. The Mahalanobis National Crop Forecast Centre operationally uses space-based observations for multiple pre-harvest forecasts of major crops (Indian Space Research Organisation [ISRO], 2023). The Digital Agriculture Mission further envisages farmer registries, season-wise digital crop surveys and a Krishi Decision Support System integrating information on crops, soils, weather and water resources (Press Information Bureau [PIB], 2024). These initiatives can create a far richer base for prediction - provided the data are reliable, representative and responsibly governed.

Why accuracy wins attention

AI models are commonly compared using statistics such as mean absolute error (MAE), root mean squared error (RMSE) and the coefficient of determination (R-squared). In simple terms, these measures ask how far predictions are from observed yields and how well the model follows variation in the data. They are useful and necessary. The danger begins when a single average score is treated as the complete verdict.

A model may achieve a low national RMSE because it performs extremely well in data-rich irrigated regions, while repeatedly overpredicting yields in drought-prone districts. Another model may have a slightly larger overall error but far smaller bias for rainfed farms and much better estimates of uncertainty. Choosing the first model merely because it tops an accuracy leaderboard can turn a technical success into a policy mistake.

Table 1. Two models, one misleading leaderboard

Criterion	Model A: accuracy champion	Model B: decision-ready
Overall MAE	2.2 q/ha	2.6 q/ha
Bias in rainfed districts	+4.5 q/ha	+1.2 q/ha
95% interval coverage	61%	91%
Explanation	Black-box output	Key drivers reported
Operational value	High headline accuracy	Safer local decisions

Note. Hypothetical illustration. MAE = mean absolute error; interval coverage is the share of observed yields falling inside the model's stated 95% prediction interval.

Seven questions beyond accuracy

1. Accurate where? Indian agriculture is exceptionally heterogeneous. A model trained mainly on Punjab and Haryana may not transfer reliably to eastern rainfed rice, hill agriculture or dryland pulses. Validation must therefore be spatial: test the system on districts, seasons and production environments that were not used for training. National accuracy should be accompanied by crop-wise, state-wise, irrigation-wise and farm-category performance. Averages can hide precisely the locations where a wrong forecast is most costly.

2. Accurate when? Agriculture does not stand still. New varieties are released, sowing dates shift, irrigation expands, markets alter input use and climate extremes create combinations not well represented in historical data. This is called data or concept drift: the relationship

learned from the past gradually changes. Models should be re-evaluated every season and stress-tested during droughts, floods and heatwaves rather than judged only in normal years.

3. How uncertain is the estimate? A point forecast of 4.7 tonnes per hectare suggests more certainty than the evidence may support. A range - for example, 4.1 to 5.2 tonnes - is often more honest and more useful. Procurement planning, crop insurance and food-balance decisions should account for the width and reliability of such intervals. A narrow but frequently wrong interval is worse than a wider, well-calibrated one. Uncertainty is not a weakness of science; hiding it is.

4. Can the model explain itself? Farmers and officials need not receive a computer-science lecture, but they should know the main reasons behind a forecast. Was the predicted shortfall driven by delayed sowing, a heat spell during grain filling, low vegetation vigour or inadequate soil moisture? Explainable AI methods can identify influential variables and critical growth stages, translating a black-box number into agronomic evidence (Mohan et al., 2025). Explanation also helps experts detect implausible behaviour - for example, a model claiming that yield increased because of a severe moisture deficit.

5. Fair to whom? Data-rich farms are easier for AI to learn. Smallholders, tenant farmers, women farmers, remote villages and plots with weak connectivity can become statistically invisible. If their crops are under-recorded or their management conditions are missing, the model may systematically serve them less well. Trust in agricultural AI depends on transparency, fairness, accountability, privacy and clarity about who receives benefits and bears risks (Gardezi et al., 2024). Local-language access and low-bandwidth delivery are therefore part of model quality, not optional extras.

6. Is it timely and actionable? A remarkably accurate harvest estimate released after procurement, irrigation or marketing decisions have been made has little operational value. Forecast timing should be matched to the decision: an early-season estimate for input and contingency planning, an in-season update for procurement and storage, and a near-harvest estimate for market and insurance operations. The message should also translate prediction into options: what action is suggested, for whom, by when and at what risk?

7. Who is accountable when it fails? No model should become an unchallengeable authority in high-stakes decisions. NITI Aayog's responsible-AI principles emphasise safety and reliability, equality, inclusion, privacy, transparency and accountability (NITI Aayog, 2021). UNESCO similarly calls for auditability, appropriate explainability and human oversight (UNESCO, 2021). If an AI estimate influences compensation, insurance or scheme eligibility, users need an audit trail, an identified responsible institution and an accessible grievance mechanism.

Figure 2. A broader performance standard for trustworthy agricultural AI.

An India-ready test for harvest-prediction systems

Before an AI forecast is scaled, its developers and public users should publish a short model card. This need not reveal proprietary code. It should state the prediction target, crops and regions covered, data period, missing-data treatment, validation design, error by important farmer groups, uncertainty method, update frequency, known limitations and responsible contact. The Food and Agriculture Organization has stressed that real-world testing, quality local data, training, connectivity and social context are essential; a technically effective tool can still fail when users cannot access or integrate it into decisions (Food and Agriculture Organization [FAO], 2025).

Table 2. A seven-part pre-deployment check

Test	Question to ask	Minimum good practice
Purpose	Which decision will the forecast improve?	Define user, timing and action before modelling
Data	Who and which regions are represented?	Document sources, gaps, consent and exclusions
Validation	Does it work outside the training sample?	Use future seasons and unseen locations
Uncertainty	How often is the stated range correct?	Report calibrated prediction intervals
Fairness	Which groups receive larger errors?	Publish subgroup errors and mitigation
Explanation	Can agronomists understand key drivers?	Provide local, agronomically plausible reasons
Governance	Who audits, updates and answers complaints?	Assign responsibility and redressal

Keep the farmer in the forecasting loop

The most robust approach is not AI versus human expertise, but AI with agronomy, statistics and farmer knowledge. Remote sensing can detect a vegetation anomaly; extension workers can determine whether it reflects drought, pest damage, harvest timing or a cloud-contaminated image. A model can identify an unusual yield pattern; farmers can explain a new variety, a local irrigation failure or a change in management that the database missed. Such feedback should be recorded and used to improve the next version. Participatory validation is particularly important in India. Farmers, state departments, crop scientists, statisticians, remote-sensing specialists, market agencies and insurers may use the same forecast for different purposes. Co-design helps identify what level of accuracy is adequate, what uncertainty can be tolerated and how results should be communicated. It also prevents an elegant model from solving the wrong problem.

The real standard of success

AI can make crop forecasting faster, more granular and more responsive. It can reveal patterns that are difficult to detect from field reports alone and can help India move from delayed estimates towards continuous agricultural intelligence. But harvest prediction is not a guessing contest in which the smallest average error wins. A trustworthy forecast should answer five plain questions: How much crop is expected? What range is plausible? Why does the model think so? For whom and where is it reliable? What decision should follow? When these answers are available, accuracy becomes meaningful. Without them, a precise number may simply automate false confidence. The future of agricultural AI should therefore be judged not by whether the algorithm can predict the harvest, but by whether people can understand, question and use that prediction to make agriculture more resilient, inclusive and informed.

A useful forecast gives the right number for the right place and time, with an honest range, a clear reason and a path to action.

References

1. Food and Agriculture Organization of the United Nations. (2025, April 2). AI can be a game-changing solution for farmers: FAO Innovation Chief. https://www.fao.org/newsroom/detail/ai-can-be-a-game-changing-solution-for-farmers--fao-innovation-chief/en
2. Gardezi, M., Joshi, B., Rizzo, D. M., Ryan, M., Prutzer, E., Brugler, S., & Dadkhah, A. (2024). Artificial intelligence in farming: Challenges and opportunities for building trust. Agronomy Journal, 116(3), 1217-1228. https://doi.org/10.1002/agj2.21353
3. Indian Space Research Organisation. (2023, April 28). Agriculture and soil. https://www.isro.gov.in/AgricultureandSoil.html
4. Mohan, R. N. V. J., Rayanoothala, P. S., & Sree, R. P. (2025). Next-gen agriculture: Integrating AI and XAI for precision crop yield predictions. Frontiers in Plant Science, 15, Article 1451607. https://doi.org/10.3389/fpls.2024.1451607
5. NITI Aayog. (2021). Responsible AI #AIForAll: Approach document for India, Part 1 - Principles for responsible AI. Government of India. https://www.niti.gov.in/sites/default/files/2021-02/Responsible-AI-22022021.pdf
6. Press Information Bureau. (2024, September 2). Cabinet approves Digital Agriculture Mission today with an outlay of Rs. 2,817 crore. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2050966
7. United Nations Educational, Scientific and Cultural Organization. (2021). Recommendation on the ethics of artificial intelligence. https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence
8. van Klompenburg, T., Kassahun, A., & Catal, C. (2020). Crop yield prediction using machine learning: A systematic literature review. Computers and Electronics in Agriculture, 177, Article 105709. https://doi.org/10.1016/j.compag.2020.105709