Indian Farmer
Volume 13, Issue 09, 2026, Pp. 447-453
Available online at: www.indianfarmer.net
ISSN: 2394-1227 (Online)

Popular Article

OPEN ACCESS

AI in the field: when your phone becomes a plant doctor

Raman Selvakumar[1] and Paresh Chaukhande[2]

[1]*Centre for Protected Cultivation Technology, ICAR–Indian Agricultural Research Institute,New Delhi*

[2]*Division of Vegetable Science, ICAR–Indian Agricultural Research Institute, New Delhi – 110012*

Corresponding author: selvakumarsingai@gmail.com

Received: 24/08/2026 *Accepted on:02/09/2026*

ABSTRACT

Artificial intelligence has moved from research laboratories into everyday vegetable cultivation, and its entry point is a device most farming families already own. Deep convolutional neural networks trained on tens of thousands of labelled leaf photographs identify crop–disease combinations with accuracies exceeding 99 per cent under test conditions, and this capability now reaches growers through smartphone applications returning a diagnosis and management advisory in Indian languages within seconds. Drones carrying multispectral cameras map water and nutrient stress days before symptoms become visible, allowing targeted rather than blanket intervention, while machine-learning models combining weather, soil and satellite records forecast yield ahead of harvest and give farmer producer organisations a marketing instrument. Reduced accuracy under field light, uneven connectivity, data ownership and the scarcity of tropical training images remain genuine constraints, each addressed here alongside a practical starting checklist for growers.

Keywords: Artificial intelligence; deep learning; plant disease diagnosis; drone remote sensing; precision agriculture; yield prediction; vegetable crops

INTRODUCTION

Walk through any vegetable market early in the morning and you will hear the same worries repeated in a dozen languages: a tomato crop lost to late blight that was noticed a week too late, a chilli field sprayed thrice for a pest that turned out to be a nutrient deficiency, an okra harvest that arrived when prices had already crashed. Behind each of these losses lies the same problem — information that came too late, or not at all. Artificial intelligence (AI) attacks precisely this problem. It does not plough, sow or spray; it sees, learns and warns. And in a country with over 120 million farm holdings and roughly one agricultural extension worker for every thousand farmers, a tireless digital assistant in every pocket may prove as important as any new variety or fertiliser.

In a nutshell

- Deep-learning systems trained on tens of thousands of leaf photographs now identify crop diseases with accuracy exceeding 99% under test conditions.
- Free smartphone apps bring this capability to any grower with a basic camera phone and deliver advisories in Indian languages.
- Drones with multispectral cameras detect water and nutrient stress days before visible symptoms appear, allowing targeted — not blanket — intervention.
- Machine-learning models that combine weather, soil and satellite data are making yield forecasting a practical planning tool for farmer producer organisations and cooperatives.

The doctor in your pocket

The breakthrough that carried AI from the laboratory to the field came from an unexpected direction: photography. In 2016, researchers at EPFL in Switzerland and Penn State University in the USA trained a deep convolutional neural network — a computer programme loosely inspired by the visual cortex of the brain — on 54,306 photographs of healthy and diseased crop leaves covering 14 crop species and 26 diseases. The trained network identified the correct crop–disease combination in 99.35% of test images it had never seen before (Mohanty et al., 2016). Two years later, a study using an even larger library of about 87,800 images spanning 58 plant–disease combinations pushed success rates to 99.5% (Ferentinos, 2018). What a plant pathologist learns in years of field experience, the network absorbed from labelled photographs in days of training.

This capability now sits inside free or low-cost smartphone applications. A grower photographs a suspicious leaf; within seconds the app returns the most likely diagnosis, a confidence estimate, and a management advisory — increasingly in regional languages and even as voice messages for those uncomfortable with reading on a screen (Figure 1). Apps of this type, such as the widely downloaded Plantix, are already used by millions of farmers across India for vegetables including tomato, chilli, brinjal, okra and cucurbits. The technology is not a substitute for a trained pathologist — a point returned to below — but as a first opinion available at midnight in a village with no plant clinic, it has no real competitor.

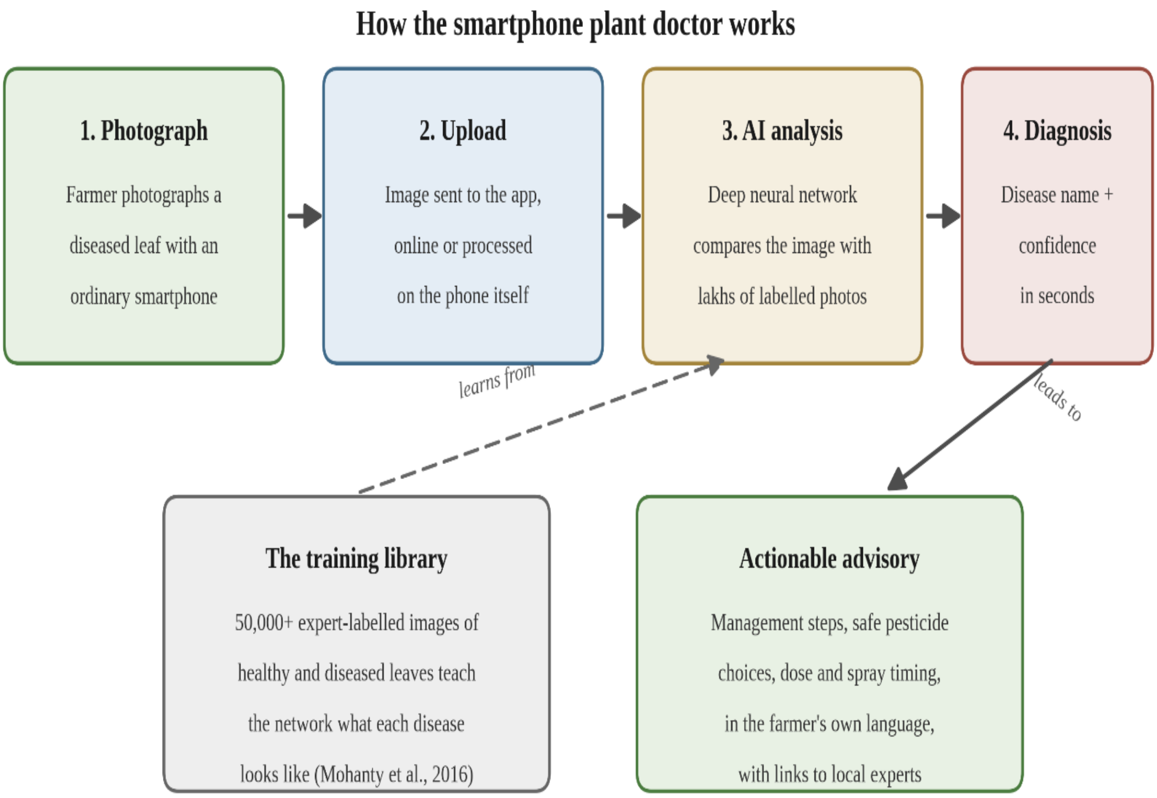

Figure 1. How the smartphone plant doctor works. A leaf photograph taken by the farmer is analysed by a deep neural network trained on tens of thousands of expert-labelled images; the diagnosis is returned in seconds together with a management advisory in the local language.

How does a machine learn to recognise a disease?

The principle is closer to apprenticeship than to programming. Nobody writes rules such as "early blight has concentric rings." Instead, the network is shown thousands of labelled examples — this image is early blight, this one is healthy, this one is a mite attack — and it adjusts millions of internal numerical connections until its own predictions match the labels. In doing so it discovers, on its own, which combinations of colour, texture, lesion shape and leaf background distinguish one problem from another. Surveys of this field report that such deep-learning approaches consistently out-perform older image-analysis techniques across agricultural tasks, from disease detection to fruit counting and weed identification (Kamilaris and Prenafeta-Boldú, 2018). The catch is equally important: the machine is only as good as its training library. A network trained on photographs taken against plain backgrounds performs noticeably worse on cluttered, sunlit field images — which is why current research effort worldwide is directed at building large, field-realistic, expert-verified image collections for each crop and region.

Eyes in the sky: drones over vegetable fields

A second family of AI tools looks at the crop not leaf by leaf but field by field. Unmanned aerial vehicles — drones — carrying multispectral cameras record how the canopy reflects light in wavelengths the human eye cannot see, particularly the near-infrared. A stressed plant alters this invisible reflectance days before it wilts or yellows visibly. Indices computed from these images, such as the widely used NDVI (normalised difference vegetation index), convert a drone flight into a colour-coded stress map of the entire field, and machine-learning models then translate those maps into specific warnings about water shortage, nitrogen deficiency or disease foci (Maes and Steppe, 2019). For a vegetable grower the practical consequence is precision: instead of spraying or fertilising the whole field on suspicion, intervention is targeted at the exact patches that need it (Figure 2). With the cost of drone services falling and custom-hiring models emerging through cooperatives and agri-startups, technology that a decade ago belonged to research stations is entering the reach of ordinary vegetable farms.

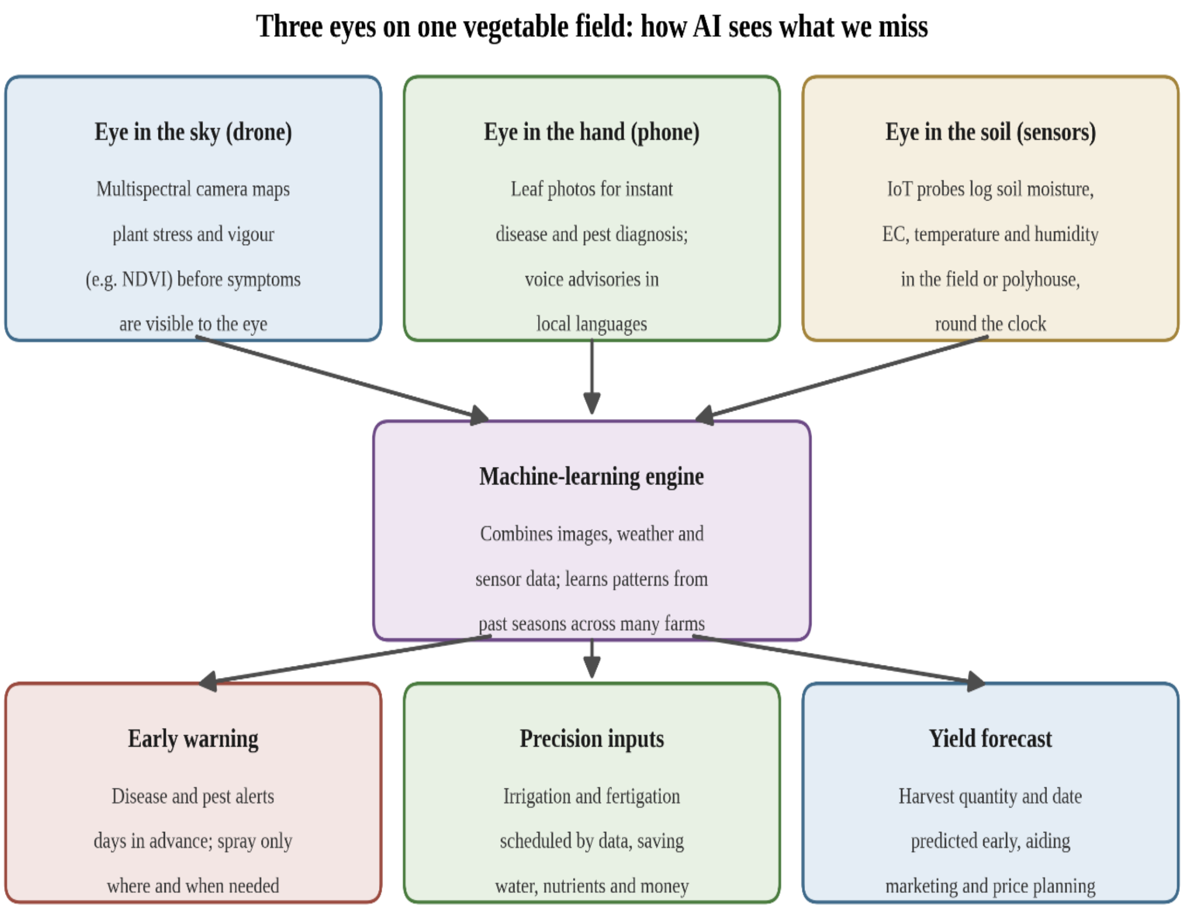

Figure 2. Three eyes on one vegetable field. Drone imagery, smartphone photographs and in-field sensors feed a machine-learning engine that returns early disease warnings, precision input scheduling and yield forecasts.

Predicting the harvest before it happens

The third capability may matter most to farm income. Machine-learning models fed with weather records, soil information, satellite imagery and past production data can forecast yield well before harvest. A systematic analysis of fifty published studies found that temperature, rainfall and soil type were the most commonly used predictors, and that neural networks were the most frequently applied and among the most accurate of the algorithms tested (van Klompenburg et al., 2020). Reviews across agriculture more broadly reach a similar conclusion: crop yield prediction is one of the areas where machine learning delivers the clearest practical gains (Liakos et al., 2018). For perishable vegetables, an early and reliable estimate of quantity and harvest date changes the marketing game entirely — farmer producer organisations can negotiate with buyers in advance, stagger plantings to avoid glut weeks, and plan cold-chain and transport capacity instead of reacting to a mountain of produce on harvest day.

Inside the smart polyhouse

Protected cultivation is where these threads knit together most tightly. A modern polyhouse already controls temperature, humidity and fertigation; adding low-cost sensors and a machine-learning layer turns it from an automated structure into a learning one. Soil-moisture and climate sensors log conditions around the clock; the software learns how the crop in that particular structure responds; and irrigation, ventilation and nutrient dosing are then scheduled from data rather than habit. The same camera-based disease models that run on phones can run on fixed cameras inside the structure, screening the crop every day without fatigue. For high-value vegetables — coloured capsicum, cherry tomato, parthenocarpic cucumber, exotic leafy greens — the combination of protected cultivation and AI-guided management is emerging worldwide as the standard against which efficiency is measured.

Table 1. What each AI tool offers a vegetable grower, and what it needs to work

Tool	**What it does**	**What the grower needs**
Disease-diagnosis app	Names the likely disease or pest from a leaf photo in seconds and suggests management	A basic smartphone with camera; many apps are free and work in Indian languages
Drone with multispectral camera	Maps stress, water and nitrogen status of the whole field before symptoms are visible	Custom-hiring service or cooperative-owned drone; a trained operator
IoT sensors + smart controller	Schedules irrigation and fertigation from real-time soil and climate data	Sensor kit and controller; most useful in polyhouses and drip-irrigated plots
Yield-forecasting model	Predicts harvest quantity and timing from weather, soil and past-season data	Usually accessed through FPOs, agri-startups or extension platforms, not individually
Voice/chat advisory assistant	Answers agronomy questions in the local language, day or night	Any phone; several government and private platforms now offer this free

What AI cannot (yet) do

Honesty about limits builds more trust than hype. First, accuracy in the field is lower than accuracy in the laboratory: unusual light, mixed infections, and diseases absent from the training library all reduce reliability, and a sensible grower treats an app's diagnosis as a strong first opinion to be confirmed — especially before any costly spray. Second, data connectivity and digital literacy remain uneven; voice interfaces and offline modes are narrowing this gap but have not closed it. Third, the models are only as unbiased as their data — a system trained largely on temperate crops will stumble on tropical vegetables and local cultivars, which is exactly why Indian institutions are building their own image libraries and region-specific models. Finally, there are questions the technology does not answer at all: who owns the data a farmer generates, and how is it protected? These are policy questions, and they deserve the same attention as the algorithms themselves.

Getting started: a five-point checklist for growers

1. Begin with a free diagnosis app on the crop you grow most; photograph clearly, in daylight, one leaf at a time.
2. Cross-check any surprising diagnosis with the local KVK, plant clinic or extension officer before purchasing chemicals.
3. For fields above one hectare, ask your FPO or cooperative about drone-based scouting on a custom-hiring basis.
4. In polyhouses, start with a single soil-moisture sensor linked to drip irrigation — the cheapest entry point into data-driven farming.
5. Keep simple written records of sowing dates, inputs and yields — your own data is the raw material every AI tool improves upon.

CONCLUSION

The direction of travel is unmistakable. Image libraries for Indian vegetable crops are growing; models are being compressed to run entirely on the phone, without internet; drone services are being organised through cooperatives and startups; and public digital-agriculture initiatives are linking soil, weather and market data into platforms that any advisory service can build upon. The next step is integration — a single assistant that sees the drone map, reads the soil sensor, recognises the leaf spot, checks the weather forecast and the mandi price, and then gives one coherent recommendation in the farmer's own language. None of the individual pieces is missing; what remains is engineering, local data and trust. The green revolution was carried into villages by seed and fertiliser. The digital one will be carried by something most farming families already own: a phone. The plant doctor is in, and the consultation fee is a photograph.

REFERENCES:

Ferentinos, K.P. (2018). Deep learning models for plant disease detection and diagnosis. Computers and Electronics in Agriculture, 145, 311–318. https://doi.org/10.1016/j.compag.2018.01.009

Kamilaris, A., & Prenafeta-Boldú, F.X. (2018). Deep learning in agriculture: A survey. Computers and Electronics in Agriculture, 147, 70–90. https://doi.org/10.1016/j.compag.2018.02.016

Liakos, K.G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. Sensors, 18(8), 2674. https://doi.org/10.3390/s18082674

Maes, W.H., & Steppe, K. (2019). Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture. Trends in Plant Science, 24(2), 152–164. https://doi.org/10.1016/j.tplants.2018.11.007

Mohanty, S.P., Hughes, D.P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. Frontiers in Plant Science, 7, 1419. https://doi.org/10.3389/fpls.2016.01419

van Klompenburg, T., Kassahun, A., & Catal, C. (2020). Crop yield prediction using machine learning: A systematic literature review. Computers and Electronics in Agriculture, 177, 105709. https://doi.org/10.1016/j.compag.2020.105709