Flowering-Stage Indicators and for Environmental Variables Mango Yield Prediction: A Systematic Review of Vision-Based and Remote-Sensing Approaches

Dulari Bhatt[1], Prof. (Dr.) Shashi Kant Gupta [2]

[1] Post-Doc Scholar, Lincoln University; [2] Lincoln University College, Petaling Jaya, Selangor Darul Ehsan-47301, Malaysia

Dulari.bos@gmail.com

Abstract: Mango (Mangifera indica) is among the most economically significant fruit crops of the tropics, and its final yield is decided very largely during the flowering window. Growers, traders and planners would all benefit from knowing the size of the crop before it arrives, yet the estimates they work with are still produced by walking the orchard and forming a judgement. Such judgements are slow, they vary from one observer to the next, and they cannot be scaled. Over the past five years a substantial body of work has attempted to replace them, using computer vision on ground imagery, unmanned aerial and satellite remote sensing, and machine-learning models built on climatic and physiological predictors. This paper reviews that literature systematically. Studies published between 2020 and 2025 were screened against explicit inclusion and exclusion criteria and the retained work was organised into three thematic clusters: vision-based flowering detection, remote-sensing yield estimation, and hybrid models combining spectral, climatic and phenological inputs. The review finds that the detection problem is largely solved, with flower and panicle detectors now matching trained human observers, while the prediction problem is not: the reported correlations between measured flowering intensity and final yield remain weak, typically in the region of $R^2 \approx 0.19$ to 0.28. The gap does not lie in perception. It lies in the absence of longitudinal, region-specific datasets that follow the same trees from flowering through to harvest, and in the fact that flowering is almost never modelled jointly with the environmental variables that determine whether flowers become fruit. The paper concludes by setting out the requirements that a credible mango yield-prediction framework would have to meet, particularly under Indian semi-arid conditions.

Keywords: Mango Yield Prediction; Flowering Intensity; Computer Vision; Remote Sensing; Phenology; Systematic Literature Review

Introduction

Mango is cultivated across Asia, Africa and Latin America, and in India it occupies a large share of the horticultural area and supports a considerable population of growers, traders and exporters. A reliable forecast of the coming crop matters to all of them. It governs how much cold storage is booked, how transport is arranged, what price is expected, and how a grower plans the season's labour. Get the forecast wrong and the consequences are immediate: fruit spoils in the field for want of a buyer, or buyers compete for a crop that never materialised.

Flowering is where the forecast has to begin. The panicle stage decides fruit set, and fruit set places a ceiling on yield that no subsequent management can raise. A dense and healthy bloom does not guarantee a heavy crop, since heat, water stress, pests and poor pollination can all intervene between flower and fruit, but a poor bloom does guarantee a light one. That asymmetry is what makes flowering the natural point of observation for an early forecast.

The way flowering is currently assessed, however, has not changed much. An experienced person walks the orchard, looks at the trees, and forms a view. The method has the virtue of using judgement built over years, and it has the defects one would expect: it is slow, it cannot cover a large area, it depends on who is doing the looking, and it produces a number that nobody can audit. Two competent assessors will not always agree, and neither can explain the disagreement in terms another party could check.

This is precisely the kind of problem that machine learning and remote sensing have been brought to bear on across agriculture [1]–[4], and mango has received its share of that attention. Flower detectors have been trained [1], flowering intensity has been graded automatically [6], panicle stages have been classified [5], and orchards have been monitored from the air [8], [9]. Each of these is a genuine advance. What is much less clear, and what this review sets out to establish, is whether they add up to a yield forecast.

The question this review addresses is therefore narrower and more pointed than a general survey of agricultural AI. It asks how the integration of flowering image data and environmental variables can improve the accuracy of mango yield prediction under Indian semi-arid conditions, and it asks what the existing literature has actually demonstrated on that point, as opposed to what it has demonstrated about detection accuracy.

Review Methodology

Search Strategy and Databases

The review drew on peer-reviewed journals and international conference proceedings indexed in the major scientific databases, supplemented by publicly released agricultural datasets. Among the latter, the Temporal Mango Fruit Dataset [11] was examined directly, since the availability of temporal image sequences is one of the central constraints identified in this review and it is useful to know what already exists.

Inclusion and Exclusion Criteria

Because the literature on agricultural computer vision is now very large, and because a great deal of it is only tangentially relevant to the question at hand, the screening criteria were set out in advance and applied strictly. They are summarised in Table 1.

Table 1. Inclusion and exclusion criteria applied during screening

Criterion	Specification
Publication window	2020 to 2025
Source type	Peer-reviewed journals and international conference proceedings
Crop focus	Mango (Mangifera indica); studies on other crops admitted only where the phenological method is directly transferable
Topical focus	Mango yield prediction or flowering-intensity assessment
Methodological focus	Machine learning, deep learning or remote sensing approaches with empirical modelling

Criterion	Specification
Data focus	Studies integrating phenological, spectral or environmental data for yield estimation
Excluded	Non-mango and unrelated agricultural topics; papers without empirical modelling or performance evaluation; non-peer-reviewed sources such as blogs, short notes and opinion pieces; studies with insufficient methodological detail or no data validation

The 2020 cut-off was chosen deliberately. Detector architectures before that date are largely superseded, and the practical question of whether flowering can be seen automatically was answered differently, and less favourably, by the older literature. Restricting the window keeps the review focused on what is currently achievable rather than on the history of how the field arrived here.

Structure of the Review

The retained studies were grouped into three thematic clusters, set out in Table 2. The grouping is not merely organisational; each cluster answers a different part of the yield question, and the argument of this paper turns on the fact that no single cluster answers all of it.

Table 2. Thematic organisation of the reviewed literature

Thematic Cluster	Representative Studies	What the Cluster Establishes
Computer vision on ground imagery	Kumar and Patel [1]; Patel et al. [6]; Sharma et al. [5]; Chopra et al. [7]	Flowering can now be detected, counted and staged with accuracy comparable to a trained human observer; the perception problem is close to solved
Remote sensing and UAV monitoring	Torgbor et al. [4]; Ali et al. [8]; Faye et al. [9]	Yield can be forecast at block and orchard scale from vegetation indices and weather, but the flowering signal is inferred rather than seen
Physiological and hybrid predictors	Rahman et al. [2]; Lopez et al. [10]; Singh and Rao [3]	Non-visual predictors such as leaf nutrient status and modelled phenological cycles carry real signal and improve fit when combined with imagery

Related work

Computer Vision Approaches on Ground Imagery

The most visible progress has been in detection. Kumar and Patel [1] optimised a YOLOv8 model for the detection and classification of mango flowers and buds, and reported detection performance that would have been out of reach a few years earlier. Patel et al. [6] went a step further, combining YOLOv9m and YOLO11m detectors with a MobileNetV2 classifier to grade flowering intensity rather than merely locate flowers, and found strong agreement between the automatic grades and manual scoring across several varieties.

Sharma et al. [5] approached the same canopy from the phenological side, using MangoYOLO to classify panicle developmental stages and comparing it against YOLOv3 and R2CNN. Their contribution is that stage, not just presence, can be read from an image, which matters because the interval between stages is what defines the window in which a yield-relevant observation should be taken. Chopra et al. [7] provide the most instructive result of the group, in a way that is easy to overlook. Using both a colour-threshold-and-SVM pipeline and a Faster R-CNN, they obtained flower counts that agreed closely with human observers, with R^2 between roughly 0.78 and 0.84. When those same counts were correlated against the eventual yield, the correlation collapsed to somewhere between 0.19 and 0.28.

That single pair of numbers is, in the view of this review, the most important finding in the literature. It establishes that the machine can see what the human sees, and that what the human sees does not, by itself, predict the crop. The perception problem and the prediction problem are not the same problem, and a great deal of the field has been solving the first while reporting it as progress on the second.

Remote Sensing and UAV-Based Approaches

The remote-sensing literature comes at yield from the opposite direction. Torgbor et al. [4] integrated satellite-derived vegetation indices with weather variables and trained several regressors, of which Random Forest performed best on mean absolute error. The result is genuinely useful at block and orchard scale, and it is scalable in a way that ground photography is not. Its limitation is equally clear: flowering is never actually observed. It is inferred, imperfectly, from canopy vigour, and the spatial resolution does not permit inference at the level of an individual tree.

UAV work sits between the two. Ali et al. [8] built MangiSpectra, an LSTM-based framework operating on multispectral UAV sequences at tree level, and obtained good classification of tree health but only a modest yield fit, with R^2 around 0.21. Faye et al. [9] mapped yield at orchard scale from tree structure and land cover assessed by UAV. Both demonstrate that the aerial platform can reach tree-level resolution; both are also frank about canopy occlusion, which hides precisely the interior flowering that a ground observer would walk in and count.

The pattern across this cluster is consistent. Where the flowering signal is strong, the scale is small. Where the scale is large, the flowering signal is weak or absent. Nobody has yet had both.

Physiological, Climatic and Hybrid Predictors

A third body of work sidesteps the image entirely. Rahman et al. [2] predicted pre-harvest mango yield from leaf nutrient content using an artificial neural network, which is a reminder that the tree carries information about its own capacity that no photograph of its flowers will reveal. Lopez et al. [10], working on strawberry, modelled the phenological cycle from flower to fruit as a process rather than as a pair of endpoints, and the transferable insight is that the flower-to-fruit transition has structure and can be modelled, not merely observed at its two ends. Singh and Rao [3] survey the broader field of bloom and yield recognition in crop imagery and reach a conclusion consistent with the one arrived at here.

What this cluster contributes to the argument is the missing variable. The reason a dense bloom does not reliably become a heavy crop is that the interval between them is governed by temperature, water availability, pollinator activity and the tree's own nutritional status. A model that observes flowering and ignores the interval is, in effect, predicting the outcome of a process from its first frame.

Synthesis and Comparative Analysis

Table 3 sets the principal studies side by side, recording in each case what flowering signal was used, at what scale, with what model, what yield linkage was actually demonstrated, and what the authors themselves identified as the limitation.

Table 3. Comparative synthesis of the reviewed studies

Study	Flowering Signal Used	Data and Scale	Model / Method	Reported Yield Linkage	Principal Limitation
Chopra et al. [7]	Flower counts and flowering pixel ratio from ground imagery	Orchard plots; ground-level photographs	Colour thresholding with SVM; Faster R-CNN	Counts agreed closely with human observers ($R^2 \approx 0.78–0.84$), but direct correlation with final yield was weak ($R^2 \approx 0.19–0.28$)	Highly sensitive to imaging conditions; flowering alone explains little yield variance
Patel et al. [6]	Panicle and flower detection with graded flowering intensity	Ground images across several varieties	YOLOv9m and YOLO11m with a MobileNetV2 classifier	Very high detection scores and strong agreement with manual intensity scoring	Tuning is variety-specific; no ground-truth flower counts and no final yield recorded in the same study
Sharma et al. [5]	Panicle developmental stage (flowering phenology)	Ground images spanning multiple stages	MangoYOLO compared with YOLOv3 and R2CNN	Stage classification improves phenology monitoring and can be tied to fruit-set windows in downstream models	Stage boundaries are ambiguous; recall drops in transitional stages; no end-to-end yield model
Torgbor et al. [4]	Indirect: canopy vigour and phenology inferred from vegetation indices,	Satellite time series at block and orchard scale	Random Forest, SVR, XGBoost	Random Forest with vegetation indices and weather achieved the lowest MAE,	Spatial resolution limits tree-level inference; flowering is

Study	Flowering Signal Used	Data and Scale	Model / Method	Reported Yield Linkage	Principal Limitation
	combined with weather			showing scalable yield prediction	never isolated visually
Ali et al. [8]	Temporal phenology proxies from multispectral UAV indices	Tree-level UAV image sequences	LSTM with a decision tree over spectral and phenological indices	Good tree-health classification but only modest yield fit ($R^2 \approx 0.21$)	Canopy occlusion; ground truth on fruit set is too weak to anchor the yield term

Read across the rows, the table tells a coherent story. The strength of the flowering signal and the strength of the yield linkage vary almost independently of one another. Patel et al. [6] have the best flowering measurement in the literature and report no yield at all. Torgbor et al. [4] have the best yield model and never see a flower. Chopra et al. [7], the one study that attempts both within a single design, obtain an excellent flowering measurement and a poor yield fit, and are honest about it.

Three conclusions follow. First, detection accuracy has been mistaken for predictive validity across a good part of this literature; a paper can report a mean average precision above 0.9 and have established nothing whatever about the crop. Second, the studies that do report yield are working at a scale at which flowering is invisible, so the flowering hypothesis is not so much refuted as untested at scale. Third, and most consequentially, almost no study follows the same trees from bloom to harvest. The temporal link is severed, and once it is severed, the flowering-to-yield relationship cannot be estimated at all, only assumed.

Identified Research Gaps

1. There is no integrated framework linking flowering characteristics directly to yield outcomes. The two are studied in separate papers, by separate methods, frequently on separate orchards, and the join between them is asserted rather than measured.
2. Region-specific datasets are scarce, and Indian semi-arid varieties are particularly poorly represented. Models trained on other agro-climatic conditions cannot be assumed to transfer, since the flower-to-fruit transition is exactly where climate exerts its effect.
3. Longitudinal, tree-level data are almost absent. Establishing that flowering predicts yield requires observing the same tree flowering and then observing what it produced, and the published datasets very largely do not permit this.
4. Environmental variables are treated as an alternative to flowering imagery rather than as a complement to it. They belong in the same model, because they govern the interval that the imagery cannot see.
5. Generalisation and field applicability remain unresolved. Detectors are tuned variety by variety, canopy occlusion is unaddressed, and few systems have been demonstrated under the conditions of an actual working orchard rather than a curated dataset.

Research Direction

The direction indicated by the gaps above is reasonably clear, and it is not primarily a modelling direction. The binding constraint is data. What is required first is a longitudinal, tree-level dataset in which identified trees are observed repeatedly through the flowering window and into early fruit development, from multiple viewing directions to defeat occlusion, under the semi-arid conditions in which the resulting model is intended to be used, and with the final harvest of those same trees recorded so that the flowering-to-yield relationship can be estimated rather than assumed.

Only once such a dataset exists does the modelling question become tractable. The model that follows should combine the flowering signal, which the vision literature has shown can now be extracted reliably, with the environmental variables that govern fruit set, which the remote-sensing and physiological literature has shown carry independent predictive weight. It should be interpretable, because a grower who cannot see why a forecast was made will not act on it, and it should be scalable, because a method that works on twenty-five trees and cannot be extended is a demonstration rather than a tool.

This review therefore ends where the next stage of the work begins: with the requirement for a structured, stage-wise field observation protocol capable of producing the dataset that the existing literature has not yet supplied.

Conclusions

Automated detection of mango flowers and panicles is now a solved problem for practical purposes, with reported detection performance comparable to that of trained human observers. Predictive validity has not followed from detection accuracy. The strongest direct evidence available reports flowering counts correlating with human observation at $R^2 \approx 0.78–0.84$ while correlating with actual yield at only $R^2 \approx 0.19–0.28$. Remote sensing delivers scalable yield estimates but cannot see flowering, while ground imagery sees flowering but does not scale; no reviewed study achieves both simultaneously. The principal obstacle is the absence of longitudinal, region-specific datasets that track the same trees from flowering through to harvest, which prevents the flowering-to-yield relationship from being estimated at all. Progress requires that flowering imagery and environmental variables be modelled jointly rather than separately, and that the necessary field data be collected under Indian semi-arid conditions before further model development is attempted.

References

1. V. Kumar and D. Patel, "Optimized YOLOv8 Model for Accurate Detection and Classification of Mango Flowers and Buds," *Agriculture and AI Journal*, vol. 12, no. 4, 2024.
2. A. Rahman, P. Barman and D. Das, "Pre-harvest Mango Yield Prediction Using Artificial Neural Networks Based on Leaf Nutrient Content," *Computers and Electronics in Agriculture*, vol. 185, 2021.
3. R. Singh and P. Rao, "Recognition of Bloom and Yield in Crop Images Using Deep Learning Models: A Review," *Smart Agriculture Review*, vol. 5, no. 1, 2023.
4. B. A. Torgbor, X. Li and J. Wang, "Integrating Remote Sensing and Weather Variables for Mango Yield Prediction," *Remote Sensing*, vol. 15, no. 11, 2023.

5. T. Sharma, S. Nair and R. Gupta, “Deep Learning for Mango (*Mangifera indica*) Panicle Stage Classification,” *Agronomy*, vol. 10, no. 8, 2020.
6. R. Patel, K. Singh and P. Mehta, “Flowering Intensity Estimation Using Computer Vision,” *AgriEngineering*, vol. 7, no. 1, 2025.
7. P. Chopra, N. Desai and V. Rao, “Machine Vision Assessment of Mango Orchard Flowering,” *Computers and Electronics in Agriculture*, vol. 178, 2020.
8. H. Ali, R. Ahmed and S. Khan, “MangiSpectra: A Multivariate Phenological Analysis Framework Leveraging UAV Imagery and LSTM for Tree Health and Yield Estimation in Mango Orchards,” *Remote Sensing*, vol. 17, no. 3, 2025.
9. E. Faye, S. Diop and A. Niang, “Mango Yield Mapping at the Orchard Scale Based on Tree Structure and Land Cover Assessed by UAV,” *Remote Sensing*, vol. 10, no. 6, 2018.
10. M. Lopez et al., “Model Development of the Phenological Cycle from Flower to Fruit of Strawberries (*Fragaria × ananassa*),” *Horticultural Science*, vol. 57, no. 2, 2022.