ARTIFICIAL INTELLIGENCE

IN FOOD CROP AGRICULTURE

Kim Kyung-jin, Attorney at Law

This book was organized and produced with artificial intelligence. It brings together materials collected from major languages around the world.

https://notebook.google.com/notebook/96fac57b-0c9e-495f-a778-dcd3daec88db

The link above opens the public NotebookLM collection of materials consulted while preparing this book. NotebookLM does not limit questions by the language of the sources. Open Settings, choose your preferred output language, and continue asking questions in that language.

Contents

Contents (continued)

CHAPTER 1

Digital Agricultural Infrastructure and AI Platforms

1.1

Digital Public Infrastructure (DPI) and AI Platforms for Agriculture

The fields of Haeje-myeon in Muan County, South Jeolla Province, are now empty ground overgrown with weeds. In a few years, glass greenhouses and a large sorting center will rise across 21.6 hectares of this land. In April 2026, the Ministry of Agriculture, Food and Rural Affairs selected a South Jeolla consortium as the preferred bidder for the National Agricultural AX Platform project. The project is valued at 254.6 billion won.

The number of farmers is falling, and those who remain are growing older. The weather swings more violently each year. Droughts and torrential rains now arrive at unfamiliar times, making it difficult to decide when to irrigate and when to harvest by relying only on the experience of feeling the soil by hand. Governments in several countries reached the same conclusion at around the same time. Neither an individual farm nor a single company can cope with this change alone, so the state must first build the foundation through which data can be exchanged.

Once the greenhouses are complete, this site will produce more than lettuce and tomatoes. Sensors on greenhouse roofs, satellites, and farmers' mobile phones will exchange data without pause. The foundation that a government lays in advance so that these data reach the places where they are needed is called digital public infrastructure, or DPI. No matter how good a letter may be, it cannot reach its recipient without an address on the envelope. Digital public infrastructure is an address system created by the state. It is a set of rules established in advance so that data from many institutions can move under the same conventions.

This arrangement is built in four layers. At the bottom, sensors inserted into the soil measure nitrogen, phosphorus, potassium, and moisture; weather stations measure wind and rain; and drones and satellites photograph the color of rice paddies. A single soil sensor can sometimes measure seven variables at once:

nitrogen, phosphorus, potassium, moisture, temperature, acidity, and electrical conductivity. In the field, it is therefore called a seven-in-one sensor. A single image taken by the European Space Agency's Sentinel satellite or the United States' Landsat satellite can contain an entire village.

At the second layer, these values travel through low-power wireless networks such as LoRaWAN or through fifth-generation mobile networks. An edge gateway, a small computer placed beside the field, filters out noise, reconciles units, and cleans the values. The third layer stores and computes the collected data. Once the values accumulate in a large repository called a data lake, models that combine many decision trees and reach an answer by majority vote, along with neural networks, read the data to detect disease and calculate yields.

Agriculture-specific large language models, or LLMs, are connected here as well, searching the data and producing an answer when a farmer asks a spoken question. The top layer consists of the screens and devices that people actually use. A voice assistant on a smartphone understands regional dialects, while irrigation valves open and close on their own.

Before these four layers can be built, another job must come first. It is like preparing the student roll before a school opens for a new term. Every farmer must be assigned a unique number, the rice paddies and fields that farmer cultivates must be drawn precisely on a satellite map, and the crops planted on that land must be recorded. The farmer registry, the land map, and the crop record must interlock before anyone can determine to whom the sensor readings and AI judgments accumulating in the upper layers belong. Without these records, even the most sophisticated model cannot issue a prescription tailored to a specific farm.

The use of data does not end there. Comparing the weight of grain harvested in autumn with the estimate produced by a model in spring reveals how accurate the prediction was. When an estimate misses the mark, investigators retrace what went wrong and feed that answer back into the model. At the end of every farming year, the model learns again, little by little. If this process stops, the model's judgments drift farther from conditions in the field with each passing year.

Rules are needed for data to move up and down these four layers. India's Digital Personal Data Protection Act, 2023 is being implemented in stages, and its core provisions on consent, processing, and the rights of data principals are scheduled to take effect in May 2027. Separate from this law, AgriStack presents a data-exchange structure built around consent managers. Statutory provisions and platform operating procedures cannot be treated as the same thing. A common vocabulary is needed as well if data are to be readable by other institutions.

AGROVOC, maintained by the Food and Agriculture Organization of the United Nations, is an agricultural thesaurus that ensures words such as rice and maize carry the same meaning across different software systems.

Before such an infrastructure existed, communities decided when to sow seed and irrigate by drawing on local experience and watching the sky. India is now adding public data to that judgment. According to Indian government figures, AgriStack had issued 101.8 million Farmer IDs by July 2026. As of November 2025, 235 million plots had been covered by the crop survey. Bharat Vistaar, launched in February 2026, had been used by more than 300,000 farmers and had processed 7.2 million queries by July of that year.

The figure of 52.8 million refers not to Bharat Vistaar users, but to farmers reached by a separate AI-based weather and monsoon advisory service. Even when Farmer IDs, field maps, and consultation records appear on the same screen, the performance figures of different programs must be read separately.

On the numbers alone, the program looks successful. Civil society groups, however, have not withdrawn their concerns. If Aadhaar, the biometric identity number issued to everyone in India, is combined in one place with land records and satellite location data, the damage caused by a leak of that bundle of information would be far greater than the damage from an isolated breach.

South Korea is only beginning. The National Agricultural AX Platform planned for the fields of Muan is to be operated from 2026 to 2030 through a special-purpose company, or SPC, funded jointly by the public and private sectors. The plan is to connect cultivation, sorting, and distribution in a single stream of data, but the system has not yet been operated in practice. Compared with India's AgriStack, which has handled real data from tens of millions of farms for more than five years, the Korean project remains at the design stage. This distance between plan and execution is the challenge now facing agricultural AI in South Korea.

These efforts are not proceeding in complete isolation from one country to the next. In April 2025, the Food and Agriculture Organization of the United Nations convened its first global dialogue on digital agriculture and artificial intelligence, where participants discussed principles for the development and governance of AI in agriculture. Digital agriculture innovation hubs in Dominica, Grenada, Ethiopia, and Morocco belonged to a separate initiative that had been under way since 2021.

The global dialogue and the innovation hubs are connected, but no official roadmap has prescribed a specific technical architecture in which raw data remain inside national borders while only summary values are sent abroad. Because each country has different data laws, communications networks, and agricultural administrations, it is more accurate to distinguish common principles from projects in the field.

Japan connects soil and weather data with pest and disease forecasts through WAGRI, a public data platform built by the Ministry of Agriculture, Forestry and Fisheries and affiliated research institutions. Zero-Agri, which connects with WAGRI, is not a water-level management device for rice paddies. It is an irrigation and liquid-fertilizer control system for greenhouse horticultural crops such as tomatoes, strawberries, and cucumbers. Water-level sensors and automatic supply valves for paddies have been tested in a separate water-management system developed by the National Agriculture and Food Research Organization.

Both forms of agriculture may be data-driven, but a nutrient-solution valve in a glass greenhouse and a sluice in a rice paddy are different pieces of equipment. Their names may appear around the same platform, yet they cannot be described as a single machine.

No matter how sophisticated a platform may be, it is useless without people who can operate it. In the World Economic Forum's 2025 survey, 68 percent of employers in agriculture, forestry, and fishing identified skills gaps in the labor market as a barrier to industry transformation. This was not a survey in which 68 percent of agricultural workers said they personally felt a skills gap, nor did it measure how prepared students at agricultural colleges were. What is needed is not a course that teaches only how to use a vast model.

Training in fields and classrooms must address why accuracy falls when the background of a leaf photograph changes, which readings should be questioned first when a sensor fails, and who should stop a decision when a model is confidently wrong.

Agricultural universities in several countries are establishing educational and research organizations that bring together computer science, agronomy, and data science. Professors and students alone cannot see every problem that arises in a rice paddy. Governments open administrative data, companies provide equipment and software, and farmers record changes in soil and crops. The quintuple helix model is a framework for explaining the relationship among universities, governments, companies, farmers, and local communities.

It describes how these five actors divide knowledge and responsibility; it is not a global tally of participating institutions.

Passing a farming prescription directly to a machine without human review remains dangerous. Several countries therefore require an agricultural extension officer to check an AI-generated answer before it reaches a farm. Within temporary authorization zones known as regulatory sandboxes, a new AI service is tested for a limited period and in a limited area before it is released formally. This is a minimum safeguard against an accident in which a model misreads a value and instructs someone to spray pesticide on the wrong field.

The question of who owns the road has yet to be resolved. If tractor manufacturers or giant cloud companies gather farm data within their own walls, the farmers who produced those data may be unable to retrieve and use their own records. The industry calls this vendor lock-in. It is like handing someone else the key and then being unable to open the door to your own house.

The burden is heavier for older farmers. In a mountain village with a weak connection, there may be no signal at all. Even the most sophisticated AI is powerless before a farmer who has to adjust a pair of glasses to read tiny letters on a screen. Whether in India or Japan, extension workers must still travel from village to village, pressing buttons on behalf of farmers and explaining the results aloud. Building these safeguards and data channels is expensive. If 254.6 billion won is being spent on a single site in Muan, one can imagine the cost of covering an entire country with dense digital public infrastructure.

Countries with limited public finances may be unable to enter the race at all, increasing the risk that the gap between data-rich and data-poor countries will widen.

The next destination of this trend is already partly visible. Instead of pulling every farm's raw data into one large central server, federated learning is being tested as a method that leaves the raw data on each farm's own computer and exchanges only what has been learned there. Experiments are proceeding in parallel to record the history of data exchanges on a blockchain so that it cannot be manipulated. If agentic AI, capable of judging and acting on its own, is added to this arrangement, a day may come when AI combines weather and soil readings, opens valves, and launches drones without a farmer watching a screen.

No one has yet determined who will be responsible when that judgment is wrong.

There are still no glass greenhouses or sensors in the fields of Muan. Only a few survey stakes stand among the weeds, marking the four corners of the greenhouses to come. When the first data emerge from this land a few years from now, will they belong to the farmer who raised the crop, or to the company that built the road along which the data travel?

Sources and References

1. MDPI. (2025). Big data analytics and machine learning for smart agriculture. MDPI. https://doi.org/10.3390/books978-3-7258-4878-2

2. Agriculture INDIA. (2026). Post budget webinar 2026: Bharat Vistaar, AI-powered digital public infrastructure for agriculture [Video]. YouTube. https://www.youtube.com/watch?v=H_Wfqw3lbhc

3. WBG Civil Society Engagement. (2026). CSPF: Examining how AI and digital agriculture create jobs in developing countries [Video]. YouTube. https://www.youtube.com/watch?v=9rRCamBUBpU

4. Committee on World Food Security. (2026). Morning session: Harnessing AI, digitalization and data governance for food security and nutrition [Video]. YouTube. https://www.youtube.com/watch?v=uQ-6ZL9UtvU

5. C-DAC. (2026). AI in agriculture: Transforming India's Krishi landscape [Video]. YouTube. https://www.youtube.com/watch?v=lV3_VnWiBpM

6. University of Agricultural Sciences, Dharwad. (2026). National conference on next-gen farming: Artificial intelligence, automation and digital transformation [Video]. YouTube. https://www.youtube.com/watch?v=BbcYxst1MFc

7. MIT College of Management, Pune. (2026). AgriTech, AI, and digital transformation in agriculture [Video]. YouTube.
https://www.youtube.com/watch?v=hEwX3Bj_M6o

8. Food and Agriculture Organization of the United Nations. (n.d.). Digital agriculture and AI: Innovation at FAO.
https://www.fao.org/innovation/digital-agriculture-and-ai-innovation/en

9. Food and Agriculture Organization of the United Nations. (n.d.). The global network of digital agriculture innovation hubs.
https://www.fao.org/in-action/global-network-digital-agriculture-innovation-hubs/en

10. Ministry of Agriculture, Food and Rural Affairs. (2026). Guidelines for the Competition to Establish the National Agricultural AX Platform SPC. Ministry of Agriculture, Food and Rural Affairs.
https://www.mafra.go.kr/bbs/home/791/594309/download.do

11. Press Information Bureau, Government of India. (2026). Bharat Vistaar and AI-based agricultural advisory: Official progress figures.
https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2289031

1.2

IoT Sensors and Edge AI Monitoring

At five in the morning, a gray rod about two fingers thick stands in the mud of a paddy bank. A green light at its tip flashes three times and goes dark. The flashes signal that it is measuring how much water seeped into the ground overnight and how low the soil temperature fell. The owner of the paddy is still under the covers when he first turns on his phone. A line graph on the screen shows the water levels and soil conditions in three paddies from the night before.

The gray rod is an Internet of Things, or IoT, sensor. Embedded in the soil, it measures how well electricity passes through the ground and converts that reading into an estimate of moisture. It is often installed alongside a nutrient sensor that measures nitrogen, phosphorus, and potassium concentrations in parts per million, a pH sensor that measures acidity, and a weather sensor that measures air temperature and humidity. A small camera may hang over the paddy as well. Every few seconds it photographs the rice leaves, checking for yellowing areas or insects on the grain.

Walking across a large field to inspect every part of it can take an entire day. A few sensors planted in each paddy can perform that work throughout the night.

These devices, which measure conditions in the soil, air, and leaves, send their signals along different routes. Over short distances, a ZigBee wireless network relays the signals. It exchanges acknowledgments that values have been received and retransmits them when a connection fails. The collected values flow into a small gateway on the paddy bank. Only after computing them does the gateway move physical equipment such as valves or sprinklers. Even then, it selects only the summarized results and uploads them to the cloud when an internet connection is available.

On broad plains, spacing may be determined by installing roughly one sensor for every 50 to 100 mu of paddy. Where the land is uneven or the soil varies widely, geographic information analysis is used in advance to determine where sensors should be placed more densely.

The question is where these numbers should be calculated. A system that sends sensor values to a distant server and waits for an answer experiences different delays depending on connection quality and the size of the data. The gap grows on days when mountains obstruct the signal or torrential rain disrupts a base station. Edge AI is therefore used to let a small computer beside the paddy reduce images, screen out anomalous values, and send only the results that are needed. Boards such as the Raspberry Pi and Jetson Nano perform this task.

Because the damage crops suffer from heat and water stress depends on the variety, growth stage, and duration of exposure, a sensor's judgment does not rest on a single temperature threshold.

A large model is difficult to run unchanged on a small computer. Engineers use quantization, which represents computed values with fewer bits; pruning, which removes connections with little influence; and knowledge distillation, which transfers the judgments of a large model to a smaller one. A 2026 TinyWeedNet study reported accuracy of 97.26 percent on the public DeepWeeds test dataset, an inference time of less than 90 milliseconds on an STM32H7 board, and a model size of about 1.815 megabytes. The result, however, came from a benchmark using public image data, not a long-term trial in the middle of a rice paddy.

Water savings and methane reductions from alternate wetting and drying must be examined in field trials separate from this image-classification model.

Photographs taken by the camera do not go straight into the model. Original images vary in size and brightness, so the edge computer first crops and resizes them into squares measuring 224 pixels on each side. It uses a calculation called bilinear interpolation to fill the gaps between pixels smoothly. Color values are normalized against a predetermined standard. Only after this preparation is a photograph placed before a model that distinguishes pests and diseases.

One experiment tested a camera-equipped four-wheel vehicle that followed crop rows. Running Canny edge detection and the Hough transform on a Raspberry Pi, the vehicle drove a 30-meter section of maize field ten times. Its mean lateral error was 6.2 centimeters on a cloudy day, 8.0 centimeters on a sunny day, and 12.5 centimeters at twilight. The paper includes no Kalman filter, twenty-per-second steering correction, rice-weeding blade, or fertilizer-rate decision function. Performance in following crop rows must be kept separate from performance in removing weeds or prescribing fertilizer, which belong to different trials.

Fields are not as tidy as textbooks. Batteries lose performance in winter, while mountains and torrential rain interrupt wireless signals. Sensors conserve power by remaining in sleep mode most of the time and activating only at scheduled moments or when a value changes sharply. By the end of 2025, the worldwide LoRaWAN deployment base, which sends small values over distances of several kilometers with little power, had surpassed 125 million devices. The figure of more than 25 percent reported by the LoRa Alliance is a compound annual growth rate, not growth in a single year.

Communications range and speed vary with transmission power, antenna height, terrain, and obstacles, so figures from one Japanese farm cannot be presented as universal specifications.

The summaries produced in this way travel to distant places as well. Equipment that moves beyond the paddy, such as drones and tractors, uses fourth-generation and fifth-generation mobile networks together with satellite links. In remote mountain fields where even these networks barely reach, older communications modules designed for text messages remain in service, carrying one short value even if they do so slowly. In fields so remote that there is no internet connection at all, the gateway opens a small local wireless server and exchanges values directly with a farmer's mobile app without going online.

When a connection is available, it uploads the cleaned values to a cloud database through a lightweight messaging protocol. These records are then used when the next model is trained.

Comparisons of cloud and edge computing speed and power consumption must use the same camera, the same model, and the same wireless network. A comparison collapses when response times, battery-consumption rates, and packet-loss rates taken from different papers are presented as if they were before-and-after results from one device. Reducing an image in the field clearly lowers the amount of data sent to the server and allows some decisions to continue when a connection fails. The size of the savings, however, must be reported as the result of a specific trial that identifies its equipment, task, and wireless conditions.

Detecting a sensor that is reporting the wrong value is difficult as well. A nutrient sensor left in the soil for months gradually loses metal from its electrodes because of salinity and chemical reactions. The screen may then report a low value, as if the soil were parched, even when there is enough water. The opposite can happen too: healthy soil may always appear wet. It is like a broken thermometer reading 39 degrees when a person has no fever. To filter out these errors, an edge computer runs several layers of checks every time a value arrives.

A box-plot method identifies readings that depart sharply from the average of recent values; a moving-average filter suppresses momentary spikes; and an interpolation method treats a physically impossible jump over a few seconds as suspicious and temporarily fills the gap with the previous value.

A 2025 study went a step further, proposing a way to identify failed sensors through cluster analysis that groups several sensor readings and examines both their similarity and their dispersion. Sensor faults fall into two broad categories. In one, values become biased and freeze on one side. In the other, they drift gradually over time. Both faults occur far more often in the harsh environment of fields and paddies than in a laboratory. It is not unusual for a paddy to be entirely submerged by a summer downpour or for frost to freeze contacts on a circuit board in winter.

The Rural Development Administration recommends calibrating pH and EC sensors every month and cleaning soil and foreign matter from their surfaces. However accurate a machine may be, it cannot replace the human hand that washes a sensor head after working in the soil.

A model trained separately on a board at each farm can send a central server only the parameters changed during training, rather than raw data containing a person's name or the location of a paddy. This is called federated learning. It has the advantage of sending less raw data beyond the farm, but it does not eliminate the risk that information will leak from the parameters or that a participating farm will submit malicious values. Nor is it time to say that compliance with India's personal data law has already been demonstrated. Federated learning is one design that can support privacy protection.

A legal basis, access controls, and security audits are still required separately.

Farms that have introduced edge AI have reported cases in which water use, communications traffic, and pesticide use fell or working time grew shorter. No meta-analysis, however, has established a range such as 25 to 50 percent as an average across crops. Results vary widely with soil, climate, equipment prices, communications costs, and the availability of repair personnel. The initial price of sensors and robots is a high barrier for small farms. This cost structure, rather than any single number, is why solar-powered boards, shared leasing, and contract services are discussed.

The resulting judgment can be delivered not only as text on a screen but also as a voice that understands local dialects. For an older farmer who finds reading difficult, one sentence from a speaker can feel more reliable than a graph.

Edge AI monitoring will not stop here. Systems are moving toward placing a thermal camera that reads heat and a hyperspectral camera that separates wavelengths invisible to the eye beside an ordinary color camera, then overlaying all of their data with satellite vegetation-index imagery on a board at the paddy bank. Combining simultaneous readings from several devices improves the accuracy of distinguishing drought damage from pest and disease damage.

The technology is expanding in another direction as well: connected to an agriculture-specific large language model, it can answer a farmer who asks in a regional dialect what is wrong with the paddy, using the sensor readings just taken to explain what must be done that day.

Only the filtered values reach the farmer's phone. One notice says that two paddies ran short of water overnight and their valves opened automatically. Another says that a reading from a third paddy was discarded as a sensor error. Only after checking these messages does the farmer pull on his boots.

Sources and References

1. Geethanjali, S. G., & Karan, N. (2025). Vision Transformer-Assisted IoT System for Smart Agriculture and Multi Crop Disease Detection. IARJSET, 12(8). https://doi.org/10.17148/IARJSET.2025.12832

2. Nezami, M., & Kumar, R. (2026). Multimodal Edge Intelligence for Crop Disease Detection and Irrigation Advisory in Precision Agriculture. International Journal of Science, Strategic Management and Technology,

2(5), 1-9.
https://doi.org/10.55041/ijsmt.v2i5.232

3. Hajirad, I. (2026). Precision and Smart Agriculture: Harnessing IoT for Enhanced Productivity and Sustainability. Intelligent Agriculture, 2(1), 84-105.
https://doi.org/10.54963/ia.v2i1.2408

4. Edge intelligence for smart agriculture: An artificial intelligence framework for precision farming and sustainable crop management. (2026).
https://doi.org/10.21474/jncs01/133

5. Hermantoro, Renjani, R. A., & Supriyanto, G. (2026). Artificial Intelligence and IoT-Based Precision Irrigation for Energy-Efficient Tropical Agriculture. International Journal of Management and Business Intelligence, 4(3), 685-698.
https://doi.org/10.59890/ijmbi.v4i3.25

6. AgroPest-Trust Edge : conception et audit d'un système frugal, explicable et conscient de l'incertitude pour la reconnaissance de ravageurs et stress biotiques sur smartphone. (2026).
https://doi.org/10.61532/rime262136

7. 農業農村情報通信環境整備 準備 . (2025). 第2回オンラインセミナー「 候 動と農業の未 ～ICT·IoT活用によるデ一タ 動型農業の重要性～」[Video]. YouTube.
https://www.youtube.com/watch?v=iqLmEFT1wZo

8. LoRa Alliance. (2026). LoRaWAN enters its next growth phase as massive IoT scales globally.
https://resources.lora-alliance.org/home/lorawan-enters-its-next-growth-phase-as-massive-iot-scales-globally

9. Kavinsky, M. (2026). LoRaWAN pushes deeper into smart agriculture with satellite-backed connectivity. IoT Business News.
https://iotbusinessnews.com/2026/07/08/lorawan-pushes-deeper-into-smart-agriculture-with-satellite-backed-connectivity/

10. Zhang, Y., Lu, Y., Martinez-Rau, L. S., Qiu, Q., & Bader, S. (2026). Real-time on-device weed identification using a hardware-efficient lightweight CNN. Frontiers in Plant Science, 17, 1747863.
https://doi.org/10.3389/fpls.2026.1747863

11. A novel data healing framework for outlier detection and recovery approach in the internet of agriculture things sensor network. (2025). Quality & Quantity.
https://doi.org/10.1007/s11135-025-02402-5

12. Alda , M. C., & Eker, B. (2026). Development and Field Evaluation of an Autonomous Four-Wheel-Drive Agricultural Vehicle Tracking Crop Rows Using Computer Vision Technology. Journal of Agricultural Production, 7(1), 13-19.
https://doi.org/10.56430/japro.1826810

1.3

Data Governance and Resource Management

On a June afternoon, with water lapping in the paddy, a farmer crouches on the bank, sets down a shovel, and takes out a phone. A blue button appears on the screen of a smart-irrigation app. It asks whether the farmer consents to the collection of location and farming information. The instant the button is pressed, the amount of water used in the paddy, the condition of the soil, and even the rate at which the rice is growing flow to a server somewhere.

Determining who owns numbers measured in a field is difficult. The European Union Code of Conduct on Agricultural Data Sharing regards farmers as data originators or source providers and recommends that farmers be able to control the conditions of access and use by contract. It should not be read as a rule that vests farmers with ownership of all data as a legal property right. Different rights and trade secrets may attach to raw sensor readings and to forecasts processed by a company. Because the code is a recommendation rather than law, actual rights vary with service contracts and the laws of each country.

In April 2026, the U.S. state of Nebraska became the first to address this question by law. When the governor signed the Agricultural Data Privacy Act, the statute declared that soil, climate, crop, and livestock data generated on a farm belong to the farm owner as long as they have not been made public or aggregated with data from several farms. A company must first enter into a written contract with a farmer before using the data. Beginning in 2027, the resale of raw data without the farmer's written consent will itself be prohibited.

The authority responsible for penalties under the law is specified as well. The Nebraska Attorney General notifies a party of a violation and provides forty-five days to cure it. If the violation remains uncorrected, a civil penalty of one thousand dollars may be imposed for each violation. The law does not give farmers a path to sue a company directly. The government serves as the watchdog instead. Yet the law has a gap. The moment data are mixed with data from several farms, they cease to belong to one farmer alone, while the company collecting them decides whether to combine them into one large pool.

A more immediate problem can arrive before the ownership dispute is settled: the company storing the data may go out of business. If a farm that has used a smart-farm service for years receives notice one day that the app is shutting down, its accumulated soil-test records, fertilizer-use history, and annual yield graphs may disappear all at once. The farm did not copy the records to its own computer or to a paper ledger. The moment a server goes dark, it becomes difficult even to ask who owns numbers that existed only on that server.

Few farmers have made a habit of downloading their own data regularly and storing separate copies, thereby exercising what is called data portability in practice. A right written into a statute is of little use unless it becomes part of ordinary behavior. Farmers who have experienced a shutdown tend to be more cautious when entrusting data to the next service. Some decide not to use an app at all.

For consent to be genuine, the decision must be made without coercion, the person must be told in advance what will happen, and there must be a way to withdraw the decision later. A screen containing only lengthy terms and one button has difficulty meeting those conditions. The European Union Data Act grants users of connected products a right to access data generated by their devices. Its general date of application is September 12, 2025. September 12, 2026 is the reference date after which connected products and related services placed on the market must be designed to make data access easy.

The date on which a right takes effect and the reference date for a product-design obligation must be read separately.

When readings from many farms accumulate on one server, the company can see a regional picture invisible to any individual farmer. India's Unified Farmer Service Interface, or UFSI, is an exchange standard designed to connect government and private services in a standard format and provide access to core registries within the scope of consent. The recorded achievement of 101.8 million AgriStack Farmer IDs by July 2026 is confirmed. Yet no official evaluation has established the combined claim that UFSI alone saved 20 to 30 percent of irrigation water while accelerating insurance payments and loan reviews.

A data-exchange function, the water-saving effect of a particular irrigation technology, and the processing time of a financial service must be evaluated separately.

Two kinds of information with different characteristics are mixed in these data. The first is personal information. The instant a landowner's name and telephone number are attached to coordinates locating a paddy, the data become personal information capable of identifying one person. The same is true of a drone photograph of a field when boundaries and lot numbers are visible. Removing names and numbers does not necessarily make the data safe. In a village with only a few paddies, neighbors may know whose field it is from the location, area, and crop type alone.

It is like removing a name tag from a piece of clothing while leaving its shape and size visible enough to guess who wears it. Calling data anonymous does not mean they are genuinely unidentifiable.

The second is trade-secret information. How much a paddy produces each year, how much fertilizer is applied, and how much is spent on labor are business information belonging to the farm. If a neighboring farm or an input supplier learns these figures, it gains an advantage over the farm when setting a price in the next transaction. Merely watching the timing and volume of fertilizer purchases can allow someone to predict what crop will be planted in the following season. Because these two forms of information differ, they should be handled differently. In practice, they are often stored together on the same server and in the same app.

The National Agricultural AX Platform promoted by the Korean government could not avoid this problem. A consortium of South Jeolla Province and Gwangju Metropolitan City was finally selected as the operator in July 2026. The plan for the 254.6 billion won platform in Muan, South Jeolla Province, came with a condition: from the beginning, it must include procedures for separately identifying and filtering personal and sensitive information in the collected data. The two categories must be distinguished from the moment of collection, not after all the data have been gathered.

If data from several sources are to be combined in an engine that automatically calculates the allocation of irrigation or fertilizer, the data must work together. The international agricultural research community often invokes the FAIR principles as the standard: data must be findable, accessible, interoperable, and reusable. The words are easy, but many obstacles arise in the field. One institution records fertilizer components in parts per million, or ppm, while another uses kilograms per hectare, or kg/ha. Even the names of administrative districts are written slightly differently from one institution to another.

If these figures are fed unchanged into one model, a model that confuses the units may calculate a fertilizer rate dramatically different from the real amount.

This is why some propose using AGROVOC, the agricultural thesaurus maintained by the Food and Agriculture Organization of the United Nations, as a common standard. Sharing data has meaning only after people agree to use the same words in the same sense. With support from the Bill & Melinda Gates Foundation, CABI, an international agricultural research organization, began a project in 2025 to support responsible sharing of agricultural data. Its conclusion was clear.

Standards can be created to organize data well from a technical standpoint, but the method for ensuring that benefits actually return to the farmers who first supplied those data must be redesigned for each place and circumstance.

In calculations that allocate water and fertilizer, paddies without sensors risk disappearing from the screen. A shortage of water on an equipped paddy appears as a number, while an unequipped paddy leaves no record even when it endures the same drought. What the Kenyan government released in March 2026 was not a final law already in force, but a draft agricultural data policy and a draft digital agriculture information bill. Public consultation took place later that month, followed by further drafting.

Rather than saying that the rights of smallholders have already been guaranteed by law, it is accurate to say that a discussion has begun about placing data ownership and access rights in legislation.

Beyond the farm gate, the picture of resource use falls out of alignment. On the farm, sensors and AI are used to save a drop of water or a handful of fertilizer. The data centers that train and operate that AI do the opposite. Their consumption of cooling water and electricity is considerable. A data center somewhere may consume as much water as the paddy saved, or more. Ways to reduce this contradiction have already been proposed: use lightweight AI, or

TinyML, which completes its computation beside a sensor instead of relying on a large, heavy model, and impose budgets that cap electricity and water consumption at data centers.

We return to the paddy bank. The farmer presses the button. The paddy's water use and soil condition remain as one line in a database somewhere. During a drought next spring, that line may become evidence used to decide whether this paddy receives water first or later. The farmer takes a photograph of the consent screen and saves it on the phone, preparing for the day when someone may ask how the data were used.

Sources and References

1. Sangwa, S., & Mutabazi, P. (2025). Artificial intelligence in agriculture: Ethical stewardship, responsible innovation, and governance for sustainable food systems. Intelligent Agriculture, 1(2), 25-45.
https://doi.org/10.54963/ia.v1i2.1684

2. Mole, M. (2025). Revolutionizing agricultural water management through AI-driven irrigation systems: A comprehensive framework for sustainable farming practices. International Journal of Computing and Engineering, 7(21), 1-11.
https://doi.org/10.47941/ijce.3092

3. Amar, S., Bori, N., & C rvers, R. (2026). From collection to control: Data governance, digital technologies, and the politics of inclusion in the digitalisation of smallholder agriculture. Outlook on Agriculture, 55(1), 19-27.
https://doi.org/10.1177/00307270251410944

4. Scobey, W. (2026). Nebraska's Agricultural Data Privacy Act and similar state efforts. National Agricultural Law Center.
https://nationalaglawcenter.org/nebraskas-agricultural-data-privacy-act-and-similar-state-efforts/

5. European Commission. (2026). Data Act. Shaping Europe's Digital Future.
https://digital-strategy.ec.europa.eu/en/policies/data-act

6. NewsPim. (2026). Private operator selected for the National Agricultural AX Platform...special-purpose company to be established within the year.
https://www.newspim.com/news/view/20260727001240

7. Penrhys-Evans, T. (2025). Making data work better in agriculture with FAIR principles. CABI Blog.
https://blog.cabi.org/2025/05/27/making-data-work-better-in-agriculture-with-fair-principles/

8. Committee on World Food Security. (2026). Harnessing AI, digitalization and data governance for food security and nutrition: Morning session [Video]. YouTube.
https://www.youtube.com/watch?v=uQ-6ZL9UtvU

9. Agriculture INDIA. (2026). Post budget webinar 2026: Bharat Vistaar, AI-powered digital public infrastructure for agriculture [Video]. YouTube.
https://www.youtube.com/watch?v=H_Wfqw3lbhc

CHAPTER 2

Crop Monitoring and Image-Based Diagnosis in Precision Agriculture

2.1

Drone (UAV) and Satellite Remote-Sensing Analysis

Early in the morning, with fog still hanging over a paddy bank, a person holding a tablet launches a drone. The aircraft flies low above the rice plants and sends back a map scattered with green, yellow, and red patches. Although it is one paddy, some areas are deep green and others have faded to yellow. To the naked eye, it is merely a uniformly green rice field.

Light creates the difference. Chlorophyll in a leaf absorbs the red light needed for photosynthesis. Near-infrared light, invisible to the human eye, strikes the layers of cells inside the leaf and bounces back. In a vigorous leaf, these cell layers are thicker and more loosely arranged, so they reflect near-infrared light more strongly. Diseased or dry leaves reflect less. The ratio between absorbed and reflected light distinguishes a healthy leaf from a wilted one. Human eyes cannot read this ratio, but sensors that capture separate wavelengths can.

The normalized difference vegetation index, or NDVI, turns this ratio into one number. It subtracts red-light reflectance from near-infrared reflectance and divides the result by the sum of the two values. One number indicates whether a field is thriving or drooping today. The index has a limitation. Once rice grows into a dense canopy and near-infrared reflectance passes a certain level, the value stops rising. It is like a scale whose dial ends at one hundred, recording the same one hundred for a person who weighs one hundred kilograms and another who weighs one hundred fifty.

The normalized difference red-edge index, or NDRE, was developed to move beyond this limit. It reads the red-edge band, a narrow range between green and near-infrared light, and distinguishes chlorophyll and nitrogen concentrations in leaves at much finer resolution.

Reflected light is not the only signal. Leaf-surface temperature changes according to how widely the stomata on the underside of the leaf open and close. A rice plant short of water closes its stomata to reduce transpiration. The leaf then warms slightly, like a person running a low fever. A thermal sensor measures this temperature change remotely and detects drought stress earlier than the human eye. A lidar sensor fires a laser tens of thousands of times and measures how long the light takes to strike the leaf canopy and return, using the result to measure crop height and volume.

A hyperspectral sensor divides a leaf into tens or hundreds of closely spaced wavelength channels, detecting even subtle color changes in cells invaded by a pathogen.

Satellites repeatedly photograph the same ground from hundreds of kilometers above. Sentinel-2 data are free and have a revisit interval of about five days. Landsat 8 and 9 are free as well. Each satellite passes the same point every sixteen days, or every eight days when the two are used together. PlanetScope observes almost daily, but it is a commercial service. Broad, free imagery is useful for viewing an area the size of a country. Imagery with shorter intervals is needed to watch a disease spread over several days. Farms must continually weigh price against temporal resolution.

South Korea added another observation satellite in 2026. In July, the Korea Aerospace Research Institute launched Compact Advanced Satellite 4 from Vandenberg Space Force Base in the United States. According to specifications released by the Ministry of Agriculture, Food and Rural Affairs, the satellite weighs 514 kilograms, has an observation swath of 120 kilometers, and provides resolution in the 5-meter class. It carries five spectral bands suited to distinguishing the growth of crops and forests, with the goal of scanning the entire Korean Peninsula once every three days.

The satellite is still undergoing performance verification and image calibration, however, so a service providing its information to farms will not begin until next year. There is a time lag between launching a satellite and learning to read its images accurately.

A drone is more like a magnifying glass trained on one corner of a satellite map. Because it flies below the clouds, it is less constrained by the cloud cover that blocks satellite imagery. Rain, fog, gusts, and low visibility still obstruct both flight and image quality. A drone can record imagery so detailed that one pixel represents only a few millimeters or centimeters, but it cannot examine a large area at once. Batteries, pilots, and flight approval are required as well.

The photographs returned by a drone do not begin as a single map. One flight produces hundreds of individual images whose edges overlap. A photogrammetric technique called structure from motion, or SfM, uses these overlapping areas as clues to stitch the hundreds of pieces together, then adds elevation data to create a three-dimensional map of the entire paddy. Another procedure must come first. A calibration panel with known reflectance must be placed at the site, and the intensity of the sunlight at that moment must be measured.

Without this reference value, photographs taken on a cloudy day and a sunny day are read against different standards. Rice that has not changed from yesterday to today may appear different.

During the monsoon, clouds may hide the ground from optical satellites for weeks. A synthetic aperture radar, or SAR, satellite can then read water and crop structure in a paddy by transmitting microwaves and measuring the returning signal. It can observe at night as well. One regional study that overlaid radar and optical imagery to classify rice-growing areas reported accuracy of about 95 percent. This is not a fixed level of performance across every country and season. Classification changes with planting time, variety, terrain, and training data.

The resulting imagery leads to work in the field. After a typhoon, running a deep-learning model on drone photographs maps where rice has lodged and how severely it has fallen, far faster and at much finer resolution than a person could count on foot. The map can be used directly in a loss assessment that determines an insurance payment. When rice false smut or blast appears, close-range drone imagery and a red-edge index are read together to distinguish diseased areas from areas short of nitrogen. The information becomes a prescription map instructing a farmer to apply more fertilizer and pesticide only in those zones.

Disease changes the reflective properties of tissue before symptoms appear on the surface. In research on Fusarium head blight in wheat, a tractor-mounted hyperspectral instrument mapped signs of infection in the field. A 2021 study in Remote Sensing used UAV hyperspectral data to extract areas affected by rice false smut. A separate 2026 study used an improved YOLOv12n model to detect diseased panicles and reported a mean accuracy of 80.7 percent. One study used wavelengths to draw diseased zones; the other located panicles in images.

In wheat fields, drones read leaf area and water stress to predict yield. In one study, the coefficient of determination varied widely by variety and growth stage, from around 0.75 to 0.99. Performance on the day young leaves were photographed was not the same as performance after heads had emerged. Models that count individual maize plants and models for targeted weeding undergo different trials as well. Plant counting, yield prediction, and herbicide savings must each be measured by a different standard.

One way to increase accuracy is to avoid insisting on a single sensor. In one study that predicted wheat yield by feeding a model data from a drone camera, a thermal sensor, and soil-moisture measurements together, prediction error fell noticeably compared with use of any one sensor. The same is true of methods that combine satellites and drones. Other researchers have fed values from a satellite's broad scan and the fine detail recorded by a drone into one machine-learning model to assess crop condition. One device fills the gaps left by another.

The way nitrogen is read is changing too. A vegetation index estimates nitrogen status indirectly from the greenness of a leaf. A hyperspectral sensor uses narrower reflectance bands, but it does not count nitrogen molecules directly. It estimates a value through a calibration model that matches laboratory measurements of nitrogen concentration with reflectance spectra. Chlorophyll, moisture, and leaf structure affect the estimate as well. Even as higher fertilizer prices increase interest, a nitrogen map on a screen remains a prediction that requires calibration and field validation.

Overlaying nitrogen-deficient zones drawn by a drone with soil data stored in a geographic information system can produce a variable-rate fertilizer prescription map. A 2025 paper in Agriculture proposed and tested a variable-rate application system combining UAV and GIS data. Producing a prescription map in a study is not the same as opening and closing the nozzles of a commercial fertilizer truck in real time. An account that extends through actual application requires a field trial specifying the applicator, position error, valve response, and control plots.

A crop follows a fixed sequence from sowing to harvest. Rice moves from the nursery to tillering, heading, and grain filling without skipping a stage. A series of vegetation-index observations taken by a satellite every few days forms a jagged wave. The values distorted by cloud cover must be distinguished from genuine changes in growth. Once a smoothing calculation has been applied, the series can be fed into a deep-learning model that remembers temporal sequences, allowing it to read the stage rice is passing through and estimate how fully the grain will ripen.

Another method continually inserts satellite measurements into long-established, process-based crop models such as APSIM and DSSAT to update the forecast. Just as a weather forecast is revised at dawn with new observations, a crop-growth model slightly revises its yield outlook every time a satellite sends a new measurement.

The readings are not always correct. On a day of low clouds, an optical satellite may mistake a shadowed zone for withered crops. If a paddy and a farm road occupy the same pixel, the computer cannot decide whether to call that square field or road. Newly emerged weeds and newly emerged maize are nearly identical in color and form, and models have been reported to count a weed as maize or maize as a weed. Accuracy often drops sharply when a model trained in one region is taken unchanged to another region with different soil color and sunlight.

In places such as the tropics or monsoon regions, where clouds and shadows are frequent, errors that divide the same crop into different crops or merge different crops into one become noticeably larger.

The speed at which this technology takes root differs by country. In a 2026 survey of women smallholder farmers in Kenya, the number who had actually used an AI-based farming tool was far smaller even among those who said they knew about the tools. The reasons overlapped: the cost of smartphone data, the burden of reading and operating a screen, and a lack of money to act on the advice. No matter how precise a map may be, it remains on the screen if there is no money or labor for the next operation.

Price is another obstacle. A properly equipped industrial drone carrying a hyperspectral camera can cost far more than an ordinary small tractor. Flight approval and personnel to process the images are needed as well. Every drone flight produces several to tens of gigabytes of photographs. Equipping the system with a small computing device to filter them immediately in the field is a heavy burden for a small farm. Instead of buying their own drones, farms in many regions are turning to a drone-as-a-service model in which a farming cooperative or service company handles photography and analysis.

Technical efforts are under way to reduce the burden. A 2023 paper in the journal Agriculture described research that combined a transformer architecture and a super-resolution technique in a deep-learning pest-detection model, shrinking the model enough to mount it directly on a drone. Instead of sending every photograph to a cloud server and waiting for an answer, the drone or a small computer in the field makes a judgment on the spot. Such a model is more useful in mountain farmland where communications fail easily or in regions with slow internet service.

In the future, a three-layer observation system is likely to take hold: a satellite first identifies an area showing an anomaly, a drone flies there for close inspection, and the result is combined with readings from sensors planted on the ground.

We return to the paddy bank. The drone lands, and one red patch remains on the screen. The image does not say whether it is disease, shadow, or merely pooled water. The final answer comes from a person in boots who walks into the paddy and touches a leaf by hand.

Sources and References

1. Integrated AI and Remote Sensing Framework for Sustainable Agriculture and Environmental Monitoring. (2026).
https://doi.org/10.64137/3108088x/ijaes-v2i2p101

2. Development of a UAV-based crop disease detection system using deep learning algorithms to enhance precision agriculture. (2026).
https://doi.org/10.2478/ijssis-2026-0031

3. Systematic Literature Review: Machine Learning and Deep Learning Approaches for Crop Yield Prediction Using Remote Sensing Data. (2026).
https://doi.org/10.37595/mediainfo.v25i2.496

4. CMAViT: Integrating Climate, Management, and Remote Sensing Data for Crop Yield Prediction With Multimodel Vision Transformers. (2026). IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing.
https://doi.org/10.1109/jstars.2026.3669511

5. Trimble Geospatial. (2026). eCognition Webinar: Deep Learning, UAVs and Precision Agriculture [Video]. YouTube.
https://www.youtube.com/watch?v=alrseDYV8Lo

6. Solvi. (2026). Webinar: From Sensors to Science: A Research Grade Drone Imaging Workflow for Agricultural Research [Video]. YouTube.
https://www.youtube.com/watch?v=xTXOXkABmHU

7. Ministry of Agriculture, Food and Rural Affairs. (2026). Dawn of South Korea's First Agricultural and Forestry Satellite Era: Establishing a Three-Day Observation Cycle for the Korean Peninsula. Ministry of Agriculture, Food and Rural Affairs Press Release.
https://www.mafra.go.kr/bbs/home/792/578337/artclView.do

8. Deep Learning Applications for Crop Mapping Using Multi-Temporal Sentinel-2 Data and Red-Edge Vegetation Indices: Integrating Convolutional and Recurrent Neural Networks. (2025). Remote Sensing, 17(18), 3207.
https://doi.org/10.3390/rs17183207

9. Yield Prediction of Winter Wheat (Triticum aestivum) Varieties Using UAV-Derived Multispectral Vegetation Indices Across Growth Stages. (2026). Agronomy, 16(9), 939.
https://doi.org/10.3390/agronomy16090939

10. AgTechNavigator. (2026). Fertiliser shock fuels interest in hyperspectral satellites for precision nitrogen use.
https://www.agtechnavigator.com/Article/2026/04/07/fertiliser-shock-fuels-interest-in-hyperspectral-satellites-for-precision-nitrogen-use/

11. Artificial intelligence (AI) adoption and its perceptions among smallholder women farmers in Kenya. (2026).
https://doi.org/10.1007/s44279-026-00626-z

2.2 Computer Vision and Growth-Phenotype (Phenomics) Analysis

A person crouching in the middle of a rice paddy holds a tape measure and a wet notebook. One rice plant is selected, and its height is measured from the base of the roots to the tip of the panicle. Moving to the next plant, the observer counts the tillers, the stems branching from one plant, by pointing to each with a finger. Even after a full day's work, one person can measure no more than a few dozen rice plants. The recorded figures are transferred into a table that evening, and the next morning the work begins again from the start.

The height, leaf shape, color, and number of panicles a plant displays are called its phenotype. A phenotype is the combined result of the genes written into a seed and the land, weather, and fertilizer that shaped its growth. This is why two seeds from the same packet can grow to different heights by autumn if one is planted in a paddy with abundant water and the other in a paddy with little. A rice plant that stands out for vigorous growth may not come from superior seed. Its paddy may merely have received more fertilizer or better irrigation. To identify a genuinely valuable gene, the same seed must therefore be planted in several paddies.

Researchers must remove the differences created by the environment and find the differences that remain. A breeder wants to distinguish what the gene produced from what the paddy's environment added. Experience and intuition once provided the eyes for finding that difference. Cameras and computation now take on the role.

Computer vision performs work beyond the capacity of human hands and eyes. A camera paired with AI photographs a plant and automatically reads its height, number of leaves, and color from the image. This method is called high-throughput phenotyping, or HTP. TraitDiscover, a system introduced in 2025, makes the combination of cameras clear. An ordinary camera recording visible color, a hyperspectral camera separating hundreds of wavelengths invisible to the human eye, a thermal camera measuring leaf-surface temperature, and a three-dimensional laser scanner measuring depth are mounted side by side on one arm.

A sensor that detects fluorescence signals indicating how much photosynthesis a leaf is actually performing is added to them. The arm sits on a three-axis drive that moves in millimeter increments, allowing it to photograph the same leaf today from precisely the same position used yesterday. Some systems turn on their own lights and continue photographing after sunset, avoiding the shadow noise created as the angle of the sun changes.

The first values can be calculated from a two-dimensional photograph. Counting how much of a top-down image is covered by green rather than soil-colored pixels gives canopy cover. Multiplying that proportion by the degree of leaf overlap produces the leaf area index, which indicates how many units of leaf area are spread over each unit of ground area. The higher the index, the larger the surface over which the rice plant can receive sunlight and conduct photosynthesis. In the past, obtaining this one value required removing the leaves one by one, placing them on a scanner, and measuring their area.

The rice plant could not resume growing after the measurement. A photograph yields the same value without breaking the plant and can be taken again the following day. A flat image, however, counts only leaves visible from above. It cannot count the layers of leaves hidden below them. A three-dimensional point cloud therefore enters next.

Cameras and lidar turn the surfaces of leaves, stems, and panicles into points that form a three-dimensional point cloud. Neural radiance fields and Gaussian splatting, both of which connect multiple images, are used to reconstruct a plant's three-dimensional structure. A preprint released in 2026 reported reducing mean reconstruction time from 6.52 minutes to 1.58 seconds across 26 plant-image sequences. It must still be described as a result from a limited test in a prepublication study, not as a field standard that has completed peer review.

Some vehicles moving through fields use cameras to find crop rows and steer along the line. The four-wheel vehicle discussed earlier recorded mean lateral errors ranging from 6.2 to 12.5 centimeters over a 30-meter section of maize field. The trial measured crop-row tracking. It did not segment rice and weeds into pixels and remove weeds with a blade. Research continues to place eyes and hands on one robot, but steering accuracy alone cannot stand in for weeding performance or the area covered in a day.

A neural network separates the stem from the point at which each leaf begins within a point cloud. In a 2025 maize study in Agriculture, segmentation accuracy was 92.6 percent on the training data and 86.1 percent on the test data. Seven traits calculated from the segmented point cloud showed high correlation with human measurements. Accuracy, however, does not reveal how many plants can be processed in a day. Before this can become field equipment that measures thousands of plants, the time required for imaging, transmission, and point-cloud reconstruction must be measured again.

When a camera photographs the same paddy from the same position every day, it creates a dated record rather than one isolated image. From changes in color and form alone, an artificial neural network distinguishes the period when young transplants take root after planting, the period when tillers branch and the plants spread, the heading period, and the period when the grain ripens. One system tested in a real paddy followed this progression with a self-steering camera and a thermal sensor, producing a daily value related to chlorophyll concentration from leaf-color analysis and an index of nitrogen status.

Farmers have long known that topdressing just before heading, during the few days when rice makes its final strong uptake of nutrients, helps the grain fill fully. Yet the exact timing of those few days has always been an estimate to the human eye. Once daily photographs accumulated, the estimate became a date. If the same method predicts when the grain will finish ripening, the date for hiring a combine can be arranged in advance.

Rice and maize are not the only crops placed before these cameras. When a hyperspectral camera and an artificial neural network are applied to images of wheat heads, they can calculate the proportion of kernels affected by Fusarium head blight without a person opening every head. A camera cannot see potato tubers underground. It can, however, measure the volume of the leaves and stems above ground and track how quickly that volume grows, providing an advance estimate of how large the underground tubers are becoming.

The cost of installing these cameras, robots, and rail systems is substantial. A rail-mounted imaging system for one greenhouse costs tens of millions of won, and adding a three-dimensional scanner and a hyperspectral camera raises the price again. The resulting photographs and point clouds rapidly consume storage, so servers and personnel to organize and compute them are required as well. Large research institutes and breeding companies can bear this cost, but it is beyond the reach of an individual farmer cultivating only a few paddies. In some countries, research funding covers the cost and places equipment at experimental stations.

Where such support does not exist, seed companies with capital are first to gain access. Most plants now standing beneath these cameras are therefore in breeding trials selecting new varieties. Ordinary farms encounter the results only after they have been turned into seed and cultivation methods.

This is what it means to read traces of genes and environment together in one photograph. Whether a leaf is dark or pale green and whether the color between its veins is even or mottled form a pattern created by the overlap between the color inherited by the variety and what the leaf experienced over the previous few days. TraitDiscover tested its reading of these patterns on soybean, maize, and rice. According to a 2026 paper reporting the results, the system detected drought stress four days before visible symptoms appeared. It identified herbicide injury twenty-four hours before a person inspecting the plant by hand recorded it.

A human eye recognizes a problem only after a leaf droops or loses color. The camera catches subtler changes in reflected light that occur earlier. Using photographs alone, the same system compared how well one variety planted in different climate zones had adapted to each local environment. With more than ten thousand breeding lines planted side by side, photographs can identify which lines tolerate drought and which are susceptible to disease without human screening. In the past, such selection took years of withholding water, transmitting disease, and watching by eye.

The camera is not always right. Wind moves the leaves, and a moving leaf appears in a different position in one photograph than in the next. The calculations that build a three-dimensional form by overlaying photographs taken from several angles assume the plant will remain still. Wind breaks that assumption. Conditions are better inside a greenhouse, where closed doors keep out the wind and lighting remains steady. The problem arises in open paddies and fields. Shadows shift every time a cloud crosses the sun.

When a leaf that was bright a moment ago suddenly enters shade, a program separating leaves from soil by color may mistake the shadow for a lesion between the veins. An experienced observer would distinguish shade from a lesion at a glance. A machine that judges only numerical values for color and brightness often guesses wrong.

In a field with dense foliage, upper leaves often hide those below. A camera cannot photograph the concealed section of a leaf and therefore calculates an area smaller than the real one. In one experiment, calculations using an occluded leaf point cloud without correction produced a mean area error of 15.88 square centimeters and a relative discrepancy of 22.11 percent from the true value. After a deep-learning correction inferred and filled the shape of the hidden leaf surface, the error fell to 6.79 square centimeters and the relative discrepancy to 8.82 percent. It did not disappear.

Accuracy also often falls sharply when a model trained under constant greenhouse lighting is moved unchanged to an open field where cloud and sun trade places several times in one day.

There is another problem, different from the moment when a model is wrong. A deep-learning neural network may label a photograph of a leaf diseased without being able to explain why. Millions of numbers mingle in its computation, forming no chain of logic that a person can follow line by line. To address this problem, researchers overlay the photograph with a red heat map showing which area the model used as evidence. If the red patch falls precisely on the real lesion, the model is more credible. If it appears over an irrelevant background, the correct answer may have been accidental.

Breeders therefore do not copy a camera's numbers without question. Faced with a suspicious result, they still turn over the leaf themselves. No breeding company will base a contract for a new variety developed over several years on an unexplained number alone.

We return to the first paddy. The hand that once crouched with a tape measure now holds a tablet. Numbers are overlaid on a photograph just sent by a rail-mounted camera. Beside it, yesterday's and the day before yesterday's readings appear in a small graph. In one square, the height reading suddenly plunges and turns red. A gust caused one leaf to hide another at the instant of the photograph. A person still decides whether to trust the number on the screen or put on boots and walk to the plant for another look.

Sources and References

1. Advancing computer vision frontiers for sustainable precision agriculture. (2025). Asian-Australasian Journal of Bioscience and Biotechnology, 10(3). https://doi.org/10.3329/aajbb.v10i3.83439

2. Chen, H., Liu, S., Wang, C., Wang, C., Gong, K., Li, Y., & Lan, Y. (2023). Point cloud completion of plant leaves under occlusion conditions based on deep learning. Plant Phenomics, 5, 0117. https://doi.org/10.34133/plantphenomics.0117

3. Development and field evaluation of an autonomous four-wheel-drive agricultural vehicle tracking crop rows using computer vision technology. (2026). https://doi.org/10.56430/japro.1826810

4. Gupta, S. K., & Agarwal, R. (2025). Comparative analysis of unmanned aerial vehicle and ground-level data using computer vision in agriculture. Cureus Journal of Computer Science, 2, es44389. https://doi.org/10.7759/s44389-025-10304-8

5. Jia, H., Zhou, W., Zhou, W., Li, Y., Lu, H., & Wu, T. (2026). The turning point of 3D plant phenotyping: 3D foundation models enable minute-to-second cross-crop reconstruction and beyond. arXiv. https://arxiv.org/abs/2607.01753

6. Pasupuleti, M. K. (2026). Vision transformers for crop intelligence and sustainable precision agriculture. International Journal of Academic and Industrial Research Innovations (IJAIRI) Research Book Series. https://doi.org/10.62311/nesx/rb-978-81-685658-6-9

7. Selvaraj, M. (2025). AI-powered phenomics: Transforming crop science for climate-resilient future [Video]. IPPN (International Plant Phenotyping Network). https://www.youtube.com/watch?v=eoTvCz3YTzo

8. Xu, L., Zhang, J., Li, S., Zhu, G., Zhao, H., Yang, J., Li, W., Nie, H., An, W., Yang, B., Hu, W., & Han, Z. (2025). TraitDiscover: An automated high-throughput platform for multimodal plant phenotyping with real-time trait analysis. Smart Agricultural Technology. https://www.sciencedirect.com/science/article/pii/S2772375525009220

9. Zhao, Y., Liu, D., Wang, Z., Yao, J., Yu, J., Wang, Z., & Jia, Z. (2025). Automated phenotyping of maize from 3D point clouds using an optimized deep learning approach. Agriculture, 15(23), 2430. https://doi.org/10.3390/agriculture15232430

10. Zong, Z., Feng, B., Wang, S., Zhang, C., Li, C., & Xu, Y. (2025). Review of Research on Recognition and Monitoring of Plant Growth Phenotype Based on Deep Learning. INMATEH Agricultural Engineering, 81-99. https://doi.org/10.35633/inmateh-77-07

2.3

Early Diagnosis of Crop Diseases, Pests, and Weeds Using Deep Learning

A person standing on a paddy bank in July bends forward and turns over a rice leaf with a fingertip. A yellow streak is spreading along its edge. The eye cannot tell whether it was caused by too little water, an insect feeding on the leaf, or a fungus entering the tissue. The person raises a phone and points it at the leaf. A disease name and a response appear on the screen.

Diseases, insects, and weeds consume 20 to 40 percent of global crop production every year, according to a range reported by the Food and Agriculture Organization of the United Nations. Plant diseases cause more than 220 billion dollars in economic losses each year, while invasive insects cause more than 70 billion dollars. The same source does not state that losses would double if chemical and biological control were removed. The scale of the losses alone is reason enough for scouting and early diagnosis.

Deep learning emerged to fill this gap. Feed in one photograph, and a multilayered artificial neural network examines color, pattern, and the texture of leaf veins on its own before producing a disease name. It does not follow a rule written in advance by a person, such as declaring disease when yellow exceeds a specified percentage. Show the machine tens of thousands of diseased leaves, and it finds for itself what to examine. Under laboratory conditions, these systems detect unfamiliar diseases such as rice blast, Fusarium head blight of wheat, and potato late blight at accuracy above 98 percent.

By that number alone, they appear better than human experts.

Several kinds of models perform different tasks within this technology. A classification model that looks at an entire leaf and assigns one disease name uses families such as ResNet or EfficientNet. Detection models in the YOLO family count individual insects scattered across one paddy image and identify their locations. Segmentation models in the UNet family draw the percentage of a leaf surface occupied by a lesion and the pixel-level boundary between a rice stem and a weed leaf. Even within deep learning, a different question calls for a different model.

The problem lies in the laboratory. PlantVillage, a dataset widely used to train deep-learning models, contains roughly 54,000 photographs in which a single leaf is placed against white paper or another uniform background and photographed head-on. The lighting is even, there are no shadows, and no weeds, soil, or other leaves overlap the subject. A model trained this way scores above 99 percent on test images taken the same way. A paddy or field, however, is not a sheet of white paper.

A 2026 paper in Frontiers in Plant Science put a number on this drop. A ResNet50 model trained on PlantVillage classified 99.73 percent of test images taken under the same conditions correctly. When the same model was applied unchanged to PlantDoc photographs taken in real fields, accuracy collapsed to 32.05 percent. It lost 67.7 percentage points on the spot. Even while producing wrong answers, the model retained mean confidence of 79.76 percent. It delivered a confident misdiagnosis without signaling that it might be wrong. The researchers used GradCAM to map the areas of each photograph that guided the judgment.

On PlantVillage photographs, the model's gaze fell on the center of the lesion. On field photographs, it scattered across soil, weeds, and shadows. The two datasets differed statistically in color intensity, sharpness of leaf edges, and even the proportion of the image occupied by the leaf.

Efforts to close the gap have continued. The same paper tested a calibration technique that lowered confidence, but the risk of misdiagnosis remained high. Adversarial learning that tried to force the two domains into alignment reduced accuracy from the 32 percent range to the 25 percent range. Even DINOv2, a newer model pretrained on unlabeled images, reached only 43.2 percent. The researchers concluded that standard calibration and lightweight adaptation methods can reduce the gap but cannot eliminate it. It is like a student who studies only photographs on white paper and then panics in a real examination hall.

This is why a model cannot be sent into a paddy on the strength of one laboratory accuracy figure.

Rare diseases present another obstacle: there are few photographs to begin with. Some appear only in one region and one season, making it difficult to collect thousands of images. Researchers therefore go beyond easy methods such as rotating, flipping, and changing the brightness of photographs. They use generative adversarial networks, or GANs, and diffusion models to create realistic artificial lesions and mix them into the training data. The method raises recognition rates for diseases with little data. It can also confuse the model further when the texture of an artificial image differs subtly from a real lesion.

Weeds are difficult to distinguish because they resemble young crops in color and shape. Research has combined hyperspectral cameras with deep learning to distinguish barnyard grass at the nursery stage. An improved RetinaNet study that detected rice and eight species of paddy weeds reported mean average precision of 94.1 percent and 24.3 frames per second. These detection results were obtained on high-performance equipment. A lightweight model installed on a Raspberry Pi must be measured again for both speed and accuracy before it can be used on a paddy bank.

One study used SegNet to divide the image into three classes: rice, background, and weed. The principal weed was Sagittaria trifolia, a species of arrowhead. Mean pixel accuracy was 92.7 percent, and frequency-weighted intersection over union was 84.4 percent. The first figure measures the share of pixels classified correctly. The second measures how much the predicted region overlaps the true region. Even on the same image, the two metrics answer different questions.

Distinguishing a drought-wilted leaf from a diseased one presents a similar problem. A leaf yellowed by lack of water and one yellowed by fungal damage look almost the same to a camera. The judgment is therefore not made from an image alone. Readings from soil-moisture and air-temperature sensors are placed beside it. If the leaf is yellow even though there is enough water in the soil, the diagnosis moves toward disease. Overlaying images and sensor values reduces misdiagnosis. An explanatory technique such as GradCAM can then mark in red the precise spot on the leaf that led the model to call it diseased.

A farmer who distrusted an answer with no visible basis can inspect the evidence on the screen.

There are moments when even a GradCAM image is insufficient. If deep learning alone is asked to determine causal relationships among a pest or disease name, weather conditions, and soil condition, it may produce a plausible but wrong answer, a phenomenon known as hallucination. A knowledge graph that organizes the physiology and outbreak conditions of pests and diseases in advance can be used alongside it. Once a vision model finds a lesion, the knowledge graph compares the temperature and humidity at that time with recent rainfall records, checking whether conditions really suited the disease.

If the two judgments conflict, the system requests another inspection instead of delivering a confirmed diagnosis.

South Korea's National Crop Pest Management System operates an image-diagnosis service that identifies disease from photographs of affected leaves uploaded by farmers. In a 2024 announcement by the Rural Development Administration, mean accuracy was above 95 percent, while human recognition accuracy was 95.3 percent. At the time, the system covered 182 pests and diseases across 31 crops. Integrated forecasting that also reads weather and cultivation conditions remains under development. There is still a distance between diagnosis from one photograph and a forecast that follows an entire season.

Large farms are not the only places where such systems are urgently needed. Asian rice farming centers on smallholders who piece together many paddies the size of a hand. The less a farm can afford to call a professional pest and disease diagnostician, the more a phone in the hand becomes its only real point of consultation. As the price of equipment falls, small paddies will be the first to use the technology.

The numbers used in crop-disease diagnosis measure different things. Accuracy is the share of all answers that are correct. Recall is the proportion of relevant cases that are not missed. F1 balances precision and recall. mAP measures detection performance across several decision thresholds. Laboratory leaf photographs, drone imagery, microscope images of spores, and mobile photographs from the field differ in difficulty. Before percentages for rice, maize, and potato models are placed beside one another, their datasets and metrics must be aligned.

Intelligent insect traps on paddy banks are being developed to photograph insects drawn to a lamp and use AI to count them by species. The Rural Development Administration has said it is applying such devices to three moth and stinkbug species. This does not support a general claim that every light trap uses YOLO for real-time classification and sends an alert through Telegram and a farm app when an economic injury threshold is crossed. Target species, stage of development, and alert method must be checked separately for each device.

Connecting a pest or disease map to targeted spraying equipment allows a nozzle to open only where treatment is needed. In a three-year University of Arkansas trial, See & Spray reduced herbicide application by 43 to 59 percent. At low sensitivity, the Palmer amaranth population also showed a risk of increasing by 280 percent each year. The 2027 equipment is designed to recognize nearly four times as many weed types as the preceding generation and young weeds as small as one-quarter inch in diameter. Any account of the chemical saved must be paired with the fact that a weed missed by the sprayer can leave seed behind.

Using less chemical can reduce both cost and environmental burden. The amount of yield preserved varies with the crop, disease, timing of infection, treatment, and control plot. Even after a diseased spot is found early, losses may not fall if the chemical is ineffective or rain arrives immediately after spraying. Equipment prices, batteries, and software-update costs influence a farm's decision as well. This is why crop-specific field trials are needed instead of one universal reduction rate.

Making a model more accurate requires photographs from more paddies. Collecting every farm's pest and disease photographs on an outside server creates privacy and security concerns. Federated learning is therefore being tested as a method that keeps photographs on each farm's device and sends only the results learned there to a central server for aggregation. The problem that each farm grows different varieties under different climates, causing accuracy to fluctuate when the results are combined, is still being addressed.

The earlier the diagnosis, the more of the harvest may be protected. A 2011 record from an Indian agricultural research institution estimated that timely application of propiconazole prevented 2 million to 3 million tonnes of losses that wheat rust could have caused. This should not be read as the measured difference in actual yield between treated and untreated regions. In a tomato leaf curl virus trial, losses of 98.43 percent were reported when infection occurred 30 days after transplanting, compared with 5.44 percent when infection occurred after 75 days. The study was limited by its nonreplicated field design.

It can illustrate the importance of timing, but the research design should not be concealed.

At sunset, the person who stood on the paddy bank at the beginning looks again at the red mark on the phone. The lesion is confined to only a few plants in one corner of the field. Instead of shouldering a sprayer, the person walks toward that corner holding a map marked with that spot alone. The rest of the paddy remains green.

Sources and References

1. Albayrak, A., Kuran, E. C., & Kayaalp, F. (2025). Deep learning approaches in precision agriculture: A comprehensive review of crop classification, disease detection, and weed detection techniques. Konya Journal of Engineering Sciences, 13(4), 1318-1360.
https://doi.org/10.36306/konjes.1583103

2. Rural Development Administration. (2024). Results of AI-Based Image Diagnosis for Crop Pests and Diseases.
https://rda.go.kr/board/boardfarminfo.do?CONTENT1=&dataNo=100000798541&mode=view&prgId=day_farmprmninfoEntry

3. Thakur, M., & Soni, B. (2026). Deep learning-based crop disease detection for precision agriculture: A survey. International Journal for Research in Applied Science & Engineering Technology, 14(3), 77562.
https://doi.org/10.22214/ijraset.2026.77562

4. Pasupuleti, M. K. (2026). Vision transformers for crop intelligence and sustainable precision agriculture. International Journal of Academic and Industrial Research Innovations Research Book Series.
https://doi.org/10.62311/nesx/rb-978-81-685658-6-9

5. Nezami, M., & Kumar, R. (2026). Multimodal Edge Intelligence for Crop Disease Detection and Irrigation Advisory in Precision Agriculture. International Journal of Science, Strategic Management and Technology, 2(5), 1-9.
https://doi.org/10.55041/ijsmt.v2i5.232

6. Ma, X., Deng, X., Qi, L., Jiang, Y., Li, H., Wang, Y., & Xing, X. (2019). Fully convolutional network for rice seedling and weed image segmentation at the seedling stage in paddy fields. PLOS ONE, 14(4), e0215676.
https://doi.org/10.1371/journal.pone.0215676

7. International Research Journal of Modernization in Engineering Technology and Science. (2025). Pest and disease detection system using machine learning for Myanmar rice crops.
https://doi.org/10.56726/irjmets81664

8. Asami, H., & Nakamura, T. (2026). Effects of an automatic weed suppression robot for paddy fields in a semi-mountainous area on weed suppression and the growth and yield of paddy rice. Japanese Journal of Crop Science, 95(1), 28-35.
https://doi.org/10.1626/jcs.95.28

9. Xiang, K., Shi, D., & Zhu, X. (2026). Quantifying the reliability gap in cross-domain plant disease classification: Benchmarking the limited efficacy of standard mitigation techniques under controlled-to-field shift. Frontiers in Plant Science, 17, Article 1826962.
https://doi.org/10.3389/fpls.2026.1826962

10. Lovett, J. (2025). Precision agriculture research measures effectiveness of See & Spray technology. University of Arkansas System Division of Agriculture, Arkansas Agricultural Experiment Station.
https://aaes.uada.edu/news/see-and-spray-research/

11. Food and Agriculture Organization of the United Nations. (n.d.). Understanding the context. FAO Pest and Pesticide Management.
https://www.fao.org/pest-and-pesticide-management/about/understanding-the-context/en/

12. Korea IT Industry News. (2026). Apple, Cactus, and Hong Kong Umbrella Tree Pest and Disease Consultations Surge in May 2026...NCPMS AI Diagnosis Enables Immediate Response.
https://www.koreaiin.com/news/903825

CHAPTER 3

Crop Growth Prediction and Resource Optimization

3.1 Machine Learning and Deep Learning Models for Yield Prediction

In late April, rice seedlings just transplanted from the seedbed rise only about two finger joints above the water in the paddy. Their roots have barely taken hold, and they have not yet shown a hint of a panicle. Yet at that same moment, a number is already displayed on the monitor in the provincial government's agricultural-administration office, estimating how many sacks of rice that field will produce in autumn. The paddy is not even fully covered in green.

Knowing the yield in advance does not mean telling fortunes. If a teacher knows how many hours a student sits at a desk each day and the score the student received on the last test, the teacher can estimate the rough range of the next score. A paddy is no different. If one gathers how much that field harvested over the past few years, how much rain fell this spring, and how much fertility remains in the soil, the outline of the autumn figure comes into view in advance. The calculation is too overwhelming to do by hand.

Each variable is easy enough to calculate, but when hundreds of thousands of paddies and fields are considered at once, it exceeds what a person can handle mentally. That is why Machine Learning and Deep Learning take over this calculation.

The data a model sees include weather, soil, satellite images, and past harvest records. Growing degree days (GDD), which measure a crop's accumulated heat, are generally calculated by subtracting the crop-specific base temperature from the average of the day's maximum and minimum temperatures, treating negative values as zero, and then adding the result. Whether to impose an upper temperature threshold depends on the model. Soil-moisture and nutrient analyses, as well as sowing, irrigation, and fertilization records, are included as well.

A satellite's Normalized Difference Vegetation Index (NDVI) indirectly shows the greenness of leaves and growth conditions, while radar data can fill gaps in optical imagery on cloudy days.

Some scholars have stacked up hundreds of papers and counted which data, and how much of each, were used. In that tally, 88 studies used temperature, the largest number, followed by 79 that used precipitation. Soil type, farm-management records, and satellite data came next, while records of pests and diseases and soil nutrient levels were used relatively less often. Putting every conceivable datum into a model does not necessarily improve the answer. Unnecessary numbers can instead cause the model to find patterns in the wrong places.

Researchers therefore devote considerable effort to selecting only the genuinely useful variables from hundreds of candidates.

There are several ways to combine data and produce an answer. Random forests aggregate the answers of many decision trees, while LSTMs remember changes over time in temperature, rainfall, and vegetation indices. Multimodal models that read weather records, farm management, and satellite images together have also emerged. The R^2 of 0.8105 reported by the 2026 FARM study was measured against labels interpolated onto 30-meter grids from county-level canola yields in Canada. A test matched against independent, actual combine-harvester measurements produced approximately R^2 0.768.

It was an evaluation on a map, one step removed from performance in predicting yield measured directly in individual fields.

Another approach combines APSIM or DSSAT, which translate crop physiology into equations, with machine learning. Adding leaf area, soil moisture, and nitrogen stress calculated by a crop model to training data allows both physiological rules and observational data to be considered together. Some studies found that hybrid models outperformed purely statistical models, but the size of the improvement varied greatly by crop, region, baseline model, and evaluation year. The mere addition of physical laws does not guarantee the same degree of error reduction every time.

When reading predictive performance, one must first distinguish the direction of the metrics. The closer the coefficient of determination, R^2, is to 1, the better it explains variation in observations; at 0, it performs at a level similar to a prediction using only the mean. RMSE and MAE measure the distance between prediction and observation, so lower values are better. To compare scores for Chinese rice, U.S. maize, and European wheat, one must check whether the same metrics and independent test data were used. Before looking at the size of a number, one must look at the basis on which it was measured.

Another obstacle is the difficulty of looking inside the model to see how it produced such a number. The more a deep-learning system has swallowed a vast body of data whole, the less visible its inner workings become. From a farmer's perspective, it is difficult to trust an unexplained number enough to plan that year's fertilization. SHAP is a tool developed to address this problem. For each individual prediction, SHAP shows in numerical terms how much each variable contributed. It can lay out item by item how much a midsummer drought lowered the expected yield and how much the clear weather that followed made up for that loss.

With such an explanation, farmers can place greater trust in the forecast.

The problem arises when a paddy encounters a summer it has never experienced before. A model trained in one region can become more inaccurate in a neighboring region with different soil, climate, varieties, and farming practices. Researchers call this domain shift. The problem becomes more pronounced under conditions rare in the training data, such as years when drought and heat waves overlap. Rather than citing an unsupported, representative figure of 95 percent for a particular place, it is better to show results from independent tests outside the training region.

Where weather stations are sparse, the weather seen by the model is sparse as well. Rainfall and temperature in mountainous areas can differ even over short distances, but when infrequent observations are spread over a broad grid, local variation disappears. Studies of Nepal's plains and mountains and Jordan's arid farmland show that the density of the observation network and data quality set the floor for prediction. The model cannot know about unrecorded rain or frost.

AgriFM, published in Frontiers in Plant Science in 2026, was trained on 2.4 million field-season time series. In retrospective evaluations using past data, it identified combined drought and heat stress 18 to 24 days earlier than a baseline using satellite data alone. The researchers stated that prospective trials are needed to demonstrate practical benefits to farmers. The figure of 18 days is not a record from a commercial service that warned fields of an approaching disaster; it is the amount of early warning achieved by reanalyzing events that had already occurred.

None of this comes free. Obtaining satellite data, purchasing weather data spanning multiple years, and renting servers to run models all cost money. A large farm cultivating thousands of hectares can readily recoup those costs through higher yields, but for a smallholder farming only a few majigi of paddies, the subscription fee itself is a burden. Governments or cooperatives therefore often purchase forecasting services on farmers' behalf and distribute them free of charge or at very low cost.

Yet accurate forecasts bring substantial gains. In the case of potatoes, knowing before digging how much would be produced that year made it possible to schedule storage reservations and shipment plans months before harvest. For farmers, that means they can plan when to hire workers and when to call trucks before autumn arrives. Farms growing wheat or soybeans can likewise learn in advance which areas are thriving and which are lagging, and reduce costs by applying additional fertilizer only to the latter areas.

For those responsible for food security, these numbers carry a different weight. If they can estimate during the summer how much the national wheat or maize harvest will fall short of a normal year, they can begin negotiating where and how much to purchase before the autumn harvest is over. Conversely, in a year expected to bring a bumper crop, there is no reason to rush to import at high prices. Even a delay of a few weeks in such decisions can mean higher prices or lost supplies in the international grain market.

Korea's agricultural and forestry satellite was launched on July 7, 2026, and successfully made its first communication contact. Its specifications of 5-meter resolution, a 120-kilometer swath, five spectral bands, and a revisit interval of about 3 days are correct. However, the satellite is still undergoing early operations and image calibration, and regular operations are scheduled to begin in 2027. There are plans to link it to agricultural monitoring and disaster assessment, but it is not the case that results have already shown automatic calculation of insurance payments or a reduction in the loss-assessment period from weeks to days.

The success of a launch and the effects of a commercial service should not be described in the same tense.

In the market, prices move the moment the numbers appear. In the United States, corn and soybean futures fluctuate whenever the Department of Agriculture releases its monthly supply-and-demand outlook. Alongside these officials' hand calculations, automated systems that use machine-learning algorithms such as XGBoost to recalculate and compare yields each day are also being tested. If a forecast becomes faster and more accurate even by a single day, the side that gets the information first gains an edge in buying and selling decisions.

Back at the April seedbed, the seedling that has risen only two finger joints above the water has not yet been seen by anyone who knows with their own eyes how many sacks it will truly become in autumn. The number inside the server has already settled on an answer, but the person crouching in the paddy and feeling the soil may sense something different. When autumn comes, which will prove right: the number on the screen or the judgment at that person's fingertips?

Sources and References

1. Chowdhury, P., Maiti, B., Mudi, L. K., & Chatterjee, S. (2026). Artificial intelligence for sustainable agriculture: A review of machine learning and deep learning techniques for crop growth and yield prediction. International Journal of Applied Resilience and Sustainability, 2(3), 176-207. https://doi.org/10.70593/deepsci.0203007

2. Haikal, M. F., & Furqon, M. A. (2026). Systematic literature review: Machine learning and deep learning approaches for crop yield prediction using remote sensing data. Media Informatika, 25(2), 126-137. https://doi.org/10.37595/mediainfo.v25i2.496

3. Jagadeesan, J., & Nagarajan, R. (2026). A review of federated machine learning for crop yield prediction: Sustainable integration of IoT and smart agriculture. Indian Journal of Computer Science and Technology, 2026(502), 532-543. https://doi.org/10.59256/indjcst.20260502059

4. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. (2026). CMAViT: Integrating climate, management, and remote sensing data for crop yield prediction with multimodal vision transformers. https://doi.org/10.1109/jstars.2026.3669511

5. Albahli, S., & Shiraz, M. (2026, July). A multimodal geospatial foundation model anticipates crop stress and yield failure across climates and species. Frontiers in Plant Science, 17, 1883297. https://doi.org/10.3389/fpls.2026.1883297

6. Nejadshamsi, S., Zhang, Y., Porth, B., Zaki, S., Porth, L., & Khoshdel, V. (2026). FARM: Crop yield prediction via regression on Prithvi's encoder for satellite sensing. AgriEngineering, 8(1), 2. https://doi.org/10.3390/agriengineering8010002

7. Ministry of Agriculture, Food and Rural Affairs. (2026, July 5). The era of Korea's first agricultural and forestry satellite begins: A launch toward science-based agricultural administration through observation of the entire Korean Peninsula every three days. Ministry of Agriculture, Food and Rural Affairs press release.
https://www.mafra.go.kr/bbs/home/792/578337/artclView.do

8. NewsPim. (2026, July 3). Korea's first agricultural and forestry satellite launched: Precise observations every three days, from crop growth to disasters.
https://www.newspim.com/news/view/20260703000799

3.2

Precision Irrigation and Variable-Rate Fertilization (VRA) Control

At midday in July, with the sun high overhead, a farmer presses the soil beside the paddy inlet with a finger. The rice plants near the point where water enters are plump and dark green, while those in a far corner of the paddy have yellow, curled leaf tips. They were planted on the same day and received the same rain, yet their appearance differs over a distance of just twenty paces. On the mobile-phone screen in the farmer's hand is a map dividing the paddy into square cells: blue cells mark areas with ample water, and red cells mark areas short of water.

Even within a single plot, the nature of the soil differs from place to place. One area has abundant sand, so water drains quickly; another has abundant clay and holds water for a long time. On sloping land, the upper part is dry and the lower part is waterlogged. Add electrical-conductivity values that measure how well electricity travels through the soil, pH values that show acidity, and a digital elevation map that reproduces contour lines, and it becomes clear that a field that looked uniform was in fact divided into several small fields.

The old approach was like handing each of thirty children in a classroom a water bottle of exactly the same size. Thirsty children and children who had already had enough received the same amount. As a result, fertilizer remained in some areas, washed away in rainwater, and flowed into streams, while other areas lacked fertilizer and produced stunted rice. Excess fertilizer leads to salinization, which makes soil salty, and flooded paddies emit still more greenhouse gases.

Precision irrigation and variable-rate fertilization are technologies that create different prescriptions for different zones of a field. Soil moisture can be measured with capacitance or TDR sensors. Electrical-conductivity sensors show the overall state of ions dissolved in soil; they do not measure nitrogen, phosphorus, and potassium separately. Ion-selective sensors must also correct for interference from moisture, temperature, and other ions. The concentration on a sensor screen is not the same thing as the nutrient deficiency a crop actually experiences.

A prescription requires looking at field sensors, soil testing, and crop growth together.

Sensors buried in the ground alone cannot reveal an entire field. Drones and satellites fill the gaps. Leaf color captured by a drone camera becomes a vegetation-index value, but once deep green foliage densely covers the field, the older index, NDVI, shows the same high value everywhere and cannot distinguish differences between zones. Switching to the newer index NDRE, which reads near-infrared wavelengths together with wavelengths slightly closer to red, separates on the map, by color, cells short of nitrogen from cells with nitrogen to spare. Satellite imagery can do the same work.

In fields that cannot afford to fly drones or bury sensors, methods that use only a few satellite images, a statistical technique called principal component analysis, and machine-learning algorithms are sometimes used instead to identify areas deficient in soil moisture and nutrients.

The cell size of a prescription map varies according to the positional accuracy of sensors, satellites, drones, and farm machinery. One meter by one meter is not the standard for every field. Spreaders linked to maps and GPS can adjust valves by zone, and irrigation valves read soil moisture together with weather forecasts. In an eight-week 2025 trial in a jalape o greenhouse, reinforcement-learning control reduced irrigation by 36.9 percent compared with conventional time-based control while maintaining soil moisture above the permanent wilting point.

This was a test of one greenhouse and one crop, and cannot be generalized as the performance of automated irrigation across all open fields.

In rice farming, alternate wetting and drying, in which paddies are filled and drained rather than kept continuously flooded, is used. In 2026, Thailand's National Science and Technology Development Agency installed HandySense, which uses ultrasonic water-level sensors, at more than 20 research sites in rice paddies in Phichit, Lampang, and Chiang Mai. The 15 to 30 percent reduction in water use and 30 to 50 percent reduction in methane are expected benefits presented by the agency, not confirmed averages measured across all those sites. Installations for data collection and field trials with verified effects must be distinguished.

Whether rice followed alternate wetting and drying can be checked with radar satellites and water-level sensors. A 2026 study reported classification accuracy of roughly 80 to 93 percent, depending on conditions. This result concerns remotely verifying irrigation implementation. It is not a field result in which a blockchain ledger immediately returned carbon-credit revenue to farmers. Using it in a carbon market requires field measurements, uncertainty calculations, independent verification, and trading rules.

For upland crops, drip irrigation and soil-moisture sensors are combined to deliver only as much water as needed around the roots. Pulse irrigation, which supplies water intermittently, can improve water-use efficiency depending on soil conditions. Research also continues on adding drone nitrogen maps and slope data to fertilization plans. Computational methods such as multi-agent systems and naïve Bayes are being tested as well, but whether they have simultaneously demonstrated water savings, reduced fertilizer use, and increased yields must be rechecked with field data for each crop.

In large fields, the flow rate of each nozzle on a center-pivot irrigation system can be varied, or satellite data can create nitrogen prescription maps. In 2026, YaraPlus in the United Kingdom offered free variable-rate nitrogen-fertilization maps to grassland farmers. The map appears free on a mobile phone, but changing fertilizer rates in the field requires a compatible spreader that can read the map. Even if the map is free, the farmer must pay for the spreader.

Yet this technology has a shadow. Sensors buried in soil for a long time can produce erratic readings as time passes because of salinity and chemical reactions, marking a place that was wet yesterday as dry today. Farms that irrigated incorrectly after trusting a wrong number sometimes turn off the prescription map and return to older methods. Every few months, someone must go out to the paddies and fields and compare sensor values with values measured by hand; in mountain-valley fields beyond wireless range, carefully installed sensors may stop without sending a single line of data.

One method for reducing such errors is to install an automatic-calibration program on a small computing device beside the sensor, filtering values immediately, but it cannot entirely replace human hands.

The cost weighs even more heavily. Equipping a farm with a variable-rate fertilization nozzle, a drone, and several sensors can cost more than the annual net income of an ordinary smallholder. Some farms buy the equipment in the first year expecting savings, then put it in storage in the second year when sensor errors and communication outages overlap, returning to shovels and visual judgment.

Lowering equipment prices alone does not solve the adoption problem. The International Food Policy Research Institute identified shortages of training, credit, supply chains, and repair personnel as barriers to small-scale irrigation technologies. Tigray is in northern Ethiopia, and the Southern Tigray Zone is the southern zone of that region. For a pump to run in the next season, there must be a route by which spare parts can arrive when equipment breaks, rules for sharing water, and income sufficient to bear debt.

The direction toward which the technology is moving includes digital twins, which reproduce farms inside computers. Trials continue in which sensor values are used to calculate soil moisture days ahead and compare multiple irrigation scenarios. The idea of using quantum computing to allocate water and electricity across farmland as a whole and trading records of reductions as carbon assets remains closer to a projection and a research proposal. Technologies whose effects have been confirmed on farms should not be placed in the same tense as computational approaches still to be tested.

Back on the paddy bank in the early morning, the red cells on the mobile-phone screen signal that the inlet should be opened wider, and the blue cells signal that it should be left alone. The farmer looks at the screen once more and places a hand on the inlet latch. What separates paddies that can use this map from those that cannot is not so much the technology's performance as whether they can afford its cost.

Sources and References

1. Nezami, M., & Kumar, R. (2026). Multimodal Edge Intelligence for Crop Disease Detection and Irrigation Advisory in Precision Agriculture. International Journal of Science, Strategic Management and Technology, 2(5), 1-9.
https://doi.org/10.55041/ijsmt.v2i5.232

2. Sheetal, V. A., Vikranth, B. M., & Adarsha Sagar, H. V. (2026). AI-driven multi-agent framework for smart irrigation and crop health monitoring in Indian rice and sugarcane farming. Plant Science Today, 13(sp2),

1-16.
https://doi.org/10.14719/pst.12797

3. Hermantoro, Rengga Arnalis Renjani, & Gani Supriyanto. (2026). Artificial intelligence and IoT-based precision irrigation for energy-efficient tropical agriculture. International Journal of Management and Business Intelligence, 4(3), 685-698.
https://doi.org/10.59890/ijmbi.v4i3.25

4. Research Square. (2026). Impact of improved irrigation technology adoption under groundwater irrigation and its impact on smallholder farmers' crop productivity, income, and food security in Southern Tigray Zone, Ethiopia (preprint).
https://doi.org/10.21203/rs.3.rs-9797920/v1

5. Agremo, & TerraXploreDrones. (2026). AI crop monitoring & variable rate applications (webinar video). YouTube.
https://www.youtube.com/watch?v=NW9g73zQ4Yo

6. Ali, A., Hussain, T., & Zahid, A. (2025). Smart irrigation technologies and prospects for enhancing water use efficiency for sustainable agriculture. AgriEngineering, 7(4), 106.
https://www.mdpi.com/2624-7402/7/4/106

7. MDPI. (2026). Alternate wetting and drying irrigated rice paddy field water status monitoring with ALOS-2 three components and IoT sensors. Remote Sensing, 18(8), 1183.
https://www.mdpi.com/2072-4292/18/8/1183

8. MDPI. (2025). Greenhouse irrigation control based on reinforcement learning. Agronomy, 15(12), 2781.
https://doi.org/10.3390/agronomy15122781

9. National Science and Technology Development Agency. (2026). HandySense alternate wetting and drying (AWD) by NSTDA advances low-carbon rice farming in Thailand. NSTDA News.
https://www.nstda.or.th/en/news/news-years-2026/handysense.html

10. Yara UK. (2026). Variable rate nitrogen mapping available free to grassland farmers. Yara UK News.
https://www.yara.co.uk/crop-nutrition/news/news/2026/variable-rate-nitrogen-mapping-available-free-to-grassland-farmers/

11. Ringler, C., & Lefore, N. (2026, June 10). Are subsidies the solution for 'costly' small-scale irrigation technologies? International Food Policy Research Institute.
https://www.ifpri.org/blog/are-subsidies-the-solution-for-costly-small-scale-irrigation-technologies/

irrigation technologies? International Food Policy Research Institute.
https://www.ifpri.org/blog/are-subsidies-the-solution-for-costly-small-scale-irrigation-technologies/

3.3 Simulation for Predicting Crop Development Stages (DVS/DVI)

One morning in June, beside a paddy where transplanting has just finished, a farmer looks at a mobile phone. A single short sentence appears on the screen. It forecasts that the rice in this paddy will head around August 15. The young seedlings moved from the nursery bed have not yet reached ankle height, yet the device in the farmer's hand has already marked a single day two months away.

That day does not arrive alone. About twenty days before heading, another application of fertilizer is needed; for several days before and after heading, the paddy water must be kept deep to prevent cold injury and high-temperature damage. Once the panicles emerge, the timing is set for applying chemicals against rice blast and false smut, and when the grains are fully mature, the date for bringing a combine into the paddy and the order for delivery to the drying facility are set as well. Predicting when rice will push out its panicles therefore means setting the entire calendar of a year's farming.

The old way of predicting this date was to accumulate and count temperature. While temperature remains below a certain threshold, rice gains almost no speed of growth; only the amount by which temperature rises above the threshold is accumulated as that day's growth. This threshold is called the base temperature, and the method of adding each day the value obtained by subtracting the base temperature from that day's mean temperature, as if putting it into a savings box, is called growing degree days (GDD). If the result of the subtraction is less than zero, zero is simply put in for that day.

Rice grows by filling this savings box without skipping a day.

The value obtained by converting accumulated temperature into a growth scale is called DVS or DVI, depending on the model. Some rice models set transplanting at 0, heading at 1, and maturity at 2. However, a rule fixing 0.2 as the seedling stage, 0.5 as the tillering stage, and 1.5 as the grain-filling stage is not common to all varieties and models. Reference points and intermediate stages are calibrated according to the crop, variety, and formula. To read a field through a single number, one must first disclose the definition used by the model that produced that number.

Rice, however, does not respond only to temperature. It also senses the shortening of day length. Rice is a plant that shifts from growing leaves and stems to making panicles only after nights lengthen, so the signal that midsummer days have begun to shorten is reflected in its growth rate as well. When the roots dry because a paddy lacks water, growth slows; conversely, rice that has experienced water shortage may rush to push out its panicles. Temperature, day length, and water together determine the day's rate of development.

The method farmers have long used was a familiar calendar. It was the rule of thumb that heading would occur a certain number of days after transplanting. But because spring and autumn temperatures fluctuate from year to year, years in which heading dates set by this method missed the actual date by a week to nearly a fortnight were not rare. If fertilizer is applied too early, rice grows excessively tall and lodges; if it is applied too late, the number of panicles falls. The price of missing one date returned as the yield of the entire paddy.

Before there was another way to calculate it, there was little to do but blame the error on the sky and wait for the next year.

Machine learning adds two major things to the accumulated-temperature calculation. One is placing side by side a model containing physical laws and a model that finds rules in records. Crop-growth models such as APSIM, DSSAT, and SIMRIW, refined for rice in Japan, express in equations how temperature, water, and nitrogen move within rice. Artificial neural networks such as long short-term memory (LSTM) networks and transformers learn patterns from past records of tens of thousands of paddies rather than equations.

When the two models are linked, the neural network fills in characteristics specific to a particular paddy that equations alone cannot explain. If the fact that rice can only advance from 0 toward 2, never turn backward, is fixed as a rule, it also prevents the neural network from making the error of calculating developmental stages in reverse.

Within one farming district, early-maturing and medium-to-late-maturing varieties are planted together. If a development index is run using only a single regional average temperature, varieties that mature early and those that mature late become mixed into one calculation. The calculation therefore separately incorporates each variety's base temperature and photoperiod sensitivity, so that even neighboring paddies trace different curves.

The other addition is correcting the calculation in real time with photographs taken from the sky. When satellites and drones photograph a paddy, they yield figures such as the leaf area index, which shows how densely leaves have grown. Comparing this figure with the state of rice inside the calculation and shifting it forward or backward only by the amount of the discrepancy is called data assimilation. It resembles a navigation system that updates the estimated arrival time whenever it receives an actual location signal.

A calculation made several months earlier cannot be expected to remain exactly right, so each photograph of the paddy corrects it little by little.

Looking inside this correction process reveals a method called the particle filter. The computer creates hundreds of copies of the same paddy and has each calculate a development index under slightly different assumptions. When a satellite image arrives, the copies that produce answers closer to that image receive greater weight, while copies that produce far-off answers are filtered out. An ensemble Kalman filter works on a similar principle, modifying a prediction slightly each time by only the difference between the observation and the calculated value.

Korea's first agricultural and forestry satellite was launched from Vandenberg Space Force Base in the United States on July 7, 2026, and successfully made its first communication contact. It has 5-meter resolution, a 120-kilometer swath, and an observation capability on a cycle of about 3 days. However, early operations and image calibration will proceed during 2026, and regular operations are scheduled to begin in 2027. It is not yet appropriate to write in the present tense that imagery of rice and soybean growth is already being delivered to farms every 3 days.

If observations at such 3-day intervals are gathered across thousands of paddies, they can forecast how much rice in a region, and eventually an entire country, will be harvested that year. The Ministry of Agriculture, Food and Rural Affairs plans to review this accumulated development information alongside supply-and-demand forecasts for other crops such as soybeans and cabbage, buying time to respond before prices rise or fall abruptly. The calculation that predicts the date when panicles emerge in one paddy thus leads to calculations that address prices at the dinner table.

A channel for explaining calculated results in terms of farmers' questions is also being prepared. The Rural Development Administration is developing an AI agent, “AI Isak,” that combines agricultural-technology materials and a language model. The planned service date stated in its official 2026 work plan is 2027. The image-based diagnosis of pests and diseases now operating on farmers' mobile phones and AI Isak, which will in the future explain heading dates and flowering periods through conversation, have different start dates.

What actually happens if this date is delayed by one day? Around 20 to 25 days before heading, there are only a few days in which to apply nitrogen fertilizer during panicle initiation, when the panicle forms inside the rice plant. Miss those days and apply it late, and the stems alone are likely to grow excessively tall and lodge; apply it too early, and the fertilizer's strength is exhausted before grains have set on the panicle, reducing the number of grains. The single figure of a development index of 0.75 determines when to open a bag of fertilizer.

The window for irrigation days is narrow as well. For several days before and after heading, paddy water must be kept deeper than usual so that panicles can form pollen properly and withstand cold spells and abnormal heat. Supply that water one day late, and the pollen has already been damaged; drain it one day early, and the paddy is dry on the very day of danger. In paddies where rice, soybeans, and wheat are planted in rotation, this one calculation determines both the day to harvest the preceding crop and the day to plant the following crop, so a one-day error also pushes back the sowing plan for the next crop.

This alert does not appear for only one paddy. Hundreds of lower paddies that share water from the same reservoir receive the calculation together. Planting dates, paddy orientation, and soil depth differ slightly from parcel to parcel, so paddies with the same variety can still show diverging development indices, and the irrigation association reorganizes the order of water delivery by the extent of that divergence.

Even after full maturity, a difference of a few days in the harvest date affects quality. Combine operations and the order at the drying facility must be planned in advance. Accumulated temperature and weather forecasts can be used to correct dates set by experience. But how much the error is reduced differs by region, variety, planting date, and evaluation year. To know what percentage of yield follows from narrowing a date by one day, one must measure typhoons, rain, and waiting time at the drying facility together.

These one or two days do not end with a single paddy. In a village where one combine is shared among several farms, if the harvest date for an earlier paddy is off, the next paddy waiting behind it, and then the one after that, are all pushed back in sequence. As development-index calculations display the scheduled dates of many paddies side by side, it has become much easier for the whole village to arrange the combine schedule in advance.

Heat waves and tropical nights unsettle the premise of the developmental calendar. In maize, high temperature and water shortage during flowering can damage pollen and fertilization. Studies of extreme heat have used KDD above 32 degrees as a heat-stress indicator for wheat and maize. Cold injury in wheat must likewise be assessed together with the variety, developmental stage, and temperature and duration of exposure. Even if the calendar date is correct, grains can remain empty if that day's weather lies outside the historical range.

The development of insects and fungi is calculated in the same way. Fall armyworm, which chews maize leaves, develops more quickly from larva to moth as temperature rises, and this rate can likewise be calculated with accumulated temperature. The problem is that the rate at which the insect develops and the rate at which maize develops move differently within the same summer. In some years, the insect grows wings a fortnight earlier than usual and encounters the crop precisely when it is at its most vulnerable developmental stage.

Recent pest and disease forecasting models therefore place the crop's development index alongside the pest's rate of development and are moving toward identifying in advance the dangerous few days when the two overlap.

An analysis published in the international journal Earth's Future in 2025 found that when drought and heat waves overlap, U.S. wheat and maize both suffer declines in yield and in the share of grain that can actually be harvested, the so-called harvestable fraction, but the extent of the decline differs according to the crop's developmental stage. In other words, fields that encountered a heat wave around flowering and fields that encountered one at an earlier stage suffered different damage despite experiencing the same heat.

A 2025 review article in an international journal that compiled growing-degree-day models for horticultural crops likewise noted that existing formulas merely add the amount by which temperature exceeds a threshold and do not account for the fact that growth can instead stop on excessively hot days. Recent models are therefore revising their calculations by separately setting aside heat on overly hot days and assigning it a penalty-like value.

Development-index calculations do not target only a day in this summer. If the multiple future-climate scenarios issued by the Intergovernmental Panel on Climate Change (IPCC) are fed into the same formula, they can also project how much earlier this paddy's heading date will occur thirty years from now. Under a scenario in which greenhouse gases continue to be emitted as they are now, the curve shows the date on which rice heads advancing a little each year; under a scenario of sharply reduced emissions, that curve is much gentler.

The more detailed the calculations become, the more problems remain. Whether rice or wheat, every variety differs in base temperature and in the extent of its response to day length, so whenever a new variety is released, it must be grown for months in a laboratory where temperature and light can be artificially controlled to obtain that variety's own values. This laboratory is called a phytotron. It is a place where temperature and day length can be adjusted at will to bring forward and observe the full process of rice growth, from the seedbed to the day panicles emerge.

But dozens of new varieties are released each year, and there is not enough time to grow all of them in one phytotron chamber. For varieties whose trials have not yet been completed, the calculation therefore borrows values from an existing variety with a similar appearance. If the two varieties' base temperatures are in fact misaligned by about half a degree, that half degree expands into an error of several days over the course of a summer.

Recently, attempts have emerged to shorten time spent in the laboratory by combining high-throughput phenotyping, which photographs hundreds of young seedlings at once to read their rate of growth, with inverse calculation methods that estimate base temperature backward from that rate. In mountain-valley paddies without a weather station, values from neighboring stations are filled in only by estimating from distance and elevation, so an error unique to that paddy remains no matter how sophisticated the calculation becomes.

Some farmers question the calculation when the date indicated by the development index differs by more than a day from the date they have long estimated by eye. Such doubt is not entirely unfounded: though rare, models have in fact missed several days when encountering an exceptionally dry spring or a late frost. Citing the inconvenience of learning mobile-phone alerts and sensor readings and the initial equipment cost, some farms have used the system for several years and then returned to the old calendar.

Digital twins are being developed in the direction of reproducing paddies and orchards inside computers and running multiple weather and irrigation scenarios in advance. In 2026, data-accumulation trials were under way in Chilean vineyards and across several crop studies, but it is not yet time to call them commercial systems that have already demonstrated reductions in actual irrigation volumes and the effects of autonomous fertilizer application. The scene in which artificial intelligence launches a drone once a development index reaches a certain value should likewise remain a future automation scenario, not a current achievement.

Back at that paddy in June, the figure of August 15 on the mobile-phone screen is a day created jointly by accumulated temperature gathered so far, satellite images arriving every 3 days, and the weather of the past several years. A person who has walked this paddy for a long time now layers their own estimate over that number. Whether that day will actually prove correct, or whether a single heat wave arriving at the end of June will shake the entire calculation, no one yet knows.

Sources and References

1. Xiao, K., Zhou, X., Gui, H., Tian, Y., Chen, X., Li, Y., Matomela, N., & Xin, Q. (2025). Drought and extreme heat reduce wheat and maize production in the United States by lowering both crop yields and harvestable fraction. Earth's Future, 13(12), e2024EF005557. https://doi.org/10.1029/2024EF005557

2. Chowdhury, P., Maiti, B., Mudi, L. K., & Chatterjee, S. (2026). Artificial intelligence for sustainable agriculture: A review of machine learning and deep learning techniques for crop growth and yield prediction. International Journal of Applied Resilience and Sustainability, 2(3), 176-207.
https://doi.org/10.70593/deepsci.0203007

3. Patil, N., Raut, K., Bhusari, S., & Gharge, A. (2025). Digital twin integration with generative AI and foundation models for real-time precision agriculture and crop resilience. International Journal of Research in Agronomy, 8(12S), 170-177.
https://doi.org/10.33545/2618060x.2025.v8.i12sc.4355

4. Kamangir, H., Sams, B. S., Dokoozlian, N., Sanchez, L., & Earles, J. M. (2026). CMAViT: Integrating climate, management, and remote sensing data for crop yield prediction with multimodal vision transformers. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 19, 8932-8945.
https://doi.org/10.1109/jstars.2026.3669511

5. Kollipara, Y. V. P., & El-Gayar, O. (2026). AI-driven crop yield prediction for climate-resilient agriculture. Conference Proceedings of the International Conference on Agriculture, Food Security and Safety, 94-106.
https://doi.org/10.32789/agrofood.2025.5107

6. Makouate, H. F., & Zude-Sasse, M. (2025). Advances in Growing Degree Days Models for Flowering to Harvest: Optimizing Crop Management with Methods of Precision Horticulture, a Review. Horticulturae, 11(12), 1415.
https://doi.org/10.3390/horticulturae11121415

7. Kalisetti, V. S., Sudharshanam, U., Venkata, N. K. M., Mandla, R., Akula, S., Bedika, M., Dharavath, B., Dogga, S., Ramakrishna Babu, A., & Javaji, C. (2025). Smart agriculture: A climate-driven approach to modelling and forecasting fall armyworm populations in maize using machine learning algorithms. Frontiers in Plant Science, 16, Article 1636412.
https://doi.org/10.3389/fpls.2025.1636412

8. Ministry of Agriculture, Food and Rural Affairs. (2026). The era of Korea's first agricultural and forestry satellite begins: A launch toward science-based agricultural administration through observation of the entire Korean Peninsula every three days. Ministry of Agriculture, Food and Rural Affairs, Office of Agricultural Industry Innovation Policy, Science and Technology Policy Division press release (2026. 7. 5.).
https://www.mafra.go.kr/bbs/home/792/578337/artclView.do

9. Nongmin Shinmun. (2025, November 19). Rural Development Administration moves to expand “AI Isak” into an agricultural-technology information service.
https://www.nongmin.com/article/20251119500634

10. FreshFruitPortal. (2026, June 17). Innovative digital twin model targets growers' real-world irrigation and traceability needs in 2026.
https://www.freshfruitportal.com/news/2026/06/17/twin-model/

CHAPTER 4

Genomics and AI-Assisted Molecular Breeding

4.1

Building Crop Genomic Big Data and Pan-Genomes

The cold-storage rooms at genetic resource centers are maintained at around 4°C. Rows of drawers are packed with paper envelopes, each labeled by hand or with a barcode. Some contain traditional rice varieties handed down for nearly a century in fields somewhere in Jeolla Province, while others hold wild rice collected from wetlands at the foot of the Himalayas. They may look much like the rice sold at a market, but the genes inside them are entirely different.

Until now, the standard way to study the rice genome has been to select just one variety from all these envelopes, decode it completely, and then read every other rice genome by aligning it against that single variety. The first variety adopted as the standard by the International Rice Genome Sequencing Project was Nipponbare, a Japanese japonica rice. For decades afterward, rice research around the world proceeded by laying out the gene sequence of this one variety like a dictionary and fitting the sequences of other varieties onto its pages.

It is like trying to catalog every word in the world with a single dictionary. When a dialect word does not appear in the standard dictionary, there is no page on which to place it. When researchers try to force genes found only in wild or traditional rice, or sequences entirely absent from other rice, into the reference genome, those sequences may be omitted altogether or attached to the wrong location. Geneticists call this phenomenon reference bias. Only the genes possessed by the standard variety are visible, while genes absent from that variety are treated as though they never existed.

The genes that disappear in this way are often the real treasures that have enabled plants to withstand drought, pests, and disease.

Removing this blind spot requires laying every edition out on the same table. Whether the crop is rice, wheat, or soybean, aligning and overlaying the genomes of hundreds or thousands of varieties reveals two layers. The layer of genes shared by all varieties is called the core genome. The layer found only in certain cultivars, landraces, or wild species is called the variable genome. The variable genome contains capabilities absent from standard varieties, such as the ability to withstand drought, take root in saline soil, or recognize unfamiliar pathogens. A map combining these two layers is called a pan-genome.

An international team of rice researchers resequenced more than 3,000 rice accessions collected around the world and released data from the 3,000 Rice Genomes Project in 2014. This dataset became a foundation for comparing millions of variants and identifying genetic differences among indica, japonica, and regional varieties. The PNAS paper that proposed the Green Super Rice concept had appeared earlier, in 2007. The chronology in which the 2007 paper emerged from the 3,000 Rice Genomes dataset is impossible.

The lineage of Green Super Rice research and the achievements of the 3,000 Rice Genomes Project should be presented side by side, without suggesting that the later data produced the earlier research.

The genes that changed the history of rice cultivation were discovered through different routes. SD1 was identified long ago as the dwarfing gene that drove the Green Revolution, while the SH1 family exists in many Asian rice varieties and is associated with seed shattering. Xa21 is a bacterial blight resistance gene discovered in wild rice. The three cannot all be grouped together as genes whose existence would have remained unknown if researchers had examined only a single reference genome.

The power of a pan-genome lies not only in showing whether a gene is present, but in revealing structural variation, presence-absence variation, varietal alleles, and regulatory regions together.

A pan-genome links the sequences shared by multiple varieties and the sequences that diverge among them into a single graph. It can place everything from single-letter differences to inversions, duplications, large insertions, and deletions on the same map. Long reads from platforms such as PacBio HiFi and Oxford Nanopore improve the quality of chromosome-scale assemblies. This does not mean, however, that the sequence of every variety must be complete from telomere to telomere to build a graph pan-genome.

Designs that combine high-quality assemblies of representative accessions with resequencing data from many accessions are also widely used.

A mung bean graph pan-genome study published in Nature Genetics in 2026 newly assembled 11 representative accessions and used 580 accessions as panels for resequencing, population genetics, and association analyses. The researchers cataloged 75,268 gene families and 66,862 structural variants and examined candidate regions associated with twenty traits through cultivation trials in five locations. This was not a study that newly assembled all 580 genomes. Distinguishing the assembly samples from the population-analysis samples reveals the map's density and evidentiary basis.

The pan-genome does not stop there. An effort to bring together and compare not only cultivated varieties and local landraces but also distantly related wild species that cannot even interbreed with them is called a super pan-genome. In the genus Oryza, in addition to rice cultivated in Asia, there is a separate rice species long grown in Africa, as well as several truly wild relatives that cannot interbreed with either of them. Adding the genomes of these distant relatives to the map brings back into view defense genes or root structures that disappeared from cultivated species long ago.

Wheat is a hexaploid crop formed by the merger of different ancestral genomes, which makes assembly difficult. A 2024 Nature study analyzed 17 Chinese varieties and cataloged 249,976 structural variants. The map revealed variants related to VRN-A1, spring and winter growth habits, and 1RS. Although it was not a map combining all the wheat in the world, it showed the breadth of structural variation even among varieties within a single country.

Soybean pan-genome research likewise assigned different roles to different samples. Researchers deeply resequenced 2,898 cultivated and wild soybean accessions and assembled 26 representative accessions among them. This cannot be described as a map in which all 2,800-plus genomes were newly assembled. Rather than presenting genes regulating water uptake and dormancy in wild soybean, drought and disease resistance, and protein content as though they were follow-up findings from a single study, the original papers that identified each trait should be cited separately.

The change is even greater in maize. The wild ancestor of modern maize is teosinte, a grass resembling a weed with only a few sparse kernels hanging from it. When the genomes of teosinte and modern maize were aligned and overlaid, several genes emerged that distinguish the transition from a thin ear bearing only a few kernels to the thick, densely packed ear seen today. Traces of the selections people made by hand in fields over thousands of years now appear as genetic coordinates on a computer screen.

The problem is that the volume of information in these accumulated maps is too great for people to read by eye. The genome of a single rice plant alone consists of an unbroken string of more than 400 million letters. A, T, G, and C are the only four letters, but their shapes alone do not reveal which arrangements cause disease and which confer the ability to withstand drought. Artificial intelligence takes on this problem. A program based on the same principle as a language model trained to read human sentences and predict the next word is applied to genetic sequences instead of sentences.

Researchers call such programs DNA language models or genomic foundation models. After being pretrained on the genomes of hundreds of plant species, a model can encounter a sequence it has never seen and use only the surrounding context to infer whether it is the start signal for a gene, a blueprint for a particular protein, or a region that will be strongly expressed. A 2026 review concluded that such genomic foundation models had reached the stage of deciphering the regulatory grammar of plants.

Genomic selection combines genetic markers in seeds with records of yield and disease resistance collected over many years to predict the performance of individuals that have not yet fully grown. Because every line need not be grown to maturity in every generation, it can shorten the time required for early selection. Yet the time needed to screen breeding lines, fix generations, conduct multi-location adaptation trials, and register a variety follows separate schedules. Even if a computer quickly selects candidates in spring, several more seasons must pass before new seed enters a farmer's storehouse in autumn.

A denser map does not immediately produce a good variety. With support from the Ministry of Agriculture, Food and Rural Affairs, the Korea Agriculture Technology Promotion Agency collected genomic information and some 4,000 molecular markers for twelve crops, including pepper and rice, and launched a digital breeding information system. Anyone can access the cloud and download the data. The next step is the problem.

An international review published in 2025 noted that no matter how much genotype data accumulates, a predictive model becomes useless without accompanying phenotype data confirming traits such as yield or disease resistance in actual fields. Reliably collecting phenotype data across many years and locations costs more than sequencing genomes. A line that grows well in one field often performs poorly in another region or another year.

Moreover, the traits that farmers await most anxiously, such as yield, are controlled in small part by dozens or hundreds of genes rather than by a single gene, and those combinations respond differently depending on the soil and weather.

After a pan-genome identifies candidate genes, crossing and gene editing follow. CRISPR excels at cutting a target sequence or changing a few letters, but targeted insertion that places a long gene and regulatory region from a wild species directly into a desired position in a cultivated species has low efficiency and imposes heavy burdens of regeneration and verification. It is not always a faster or easier route than crossing. The effects of a single edited gene on yield, quality, and other forms of resistance must be tested again in multiple environments.

The concentration of genetic data and patents in a handful of companies is a real issue. The 71 percent reported by the USDA, however, does not represent all genetic-resource patents worldwide. It is the share held by Bayer, Corteva, and Syngenta among utility patents for new varieties registered in the United States from 1976 through 2021. Patents covering methods, GM traits, and gene sequences were excluded from this tally. Monsanto was acquired by Bayer in 2018. Any discussion of patent concentration must specify the country, period, patent category, and corporate relationships.

We return to the cold-storage room. Every envelope in those drawers remains a genetic resource that must be collected by human hands, planted by human hands, and examined. No matter how precise the pan-genome map on the screen becomes, determining how fully the gene it identifies fills rice grains in an actual paddy requires waiting through a season. Among the tens of thousands of envelopes in those drawers, how many genes have yet even to be named?

Sources and References

1. Farooq, M. A., Gao, S., Hassan, M. A., et al. (2024). Artificial intelligence in plant breeding. Trends in Genetics, 40(10), 891-908.
https://www.cell.com/trends/genetics/fulltext/S0168-9525(24)00167-7

2. Jiao, C., Xie, X., Hao, C., et al. (2025). Pan-genome bridges wheat structural variations with habitat and breeding. Nature, 637(8045), 384-393.
https://doi.org/10.1038/s41586-024-08277-0

3. Shang, L., Li, X., He, H., et al. (2022). A super pan-genomic landscape of rice. Cell Research, 32(10), 878-896.
https://doi.org/10.1038/s41422-022-00685-z

4. Chen, H., et al. (2026). Graph-based pan-genome reveals structural variations associated with agronomic traits in mung bean. Nature Genetics.
https://doi.org/10.1038/s41588-026-02644-5

5. Zhu, W., Li, W., Zhang, H., & Li, L. (2025). Big data and artificial intelligence-aided crop breeding: Progress and prospects. Journal of Integrative Plant Biology, 67(3), 722-739.
https://doi.org/10.1111/jipb.13791

6. Fu, J., Zheng, S., Fan, L., Zheng, X., & Qian, Q. (2025). Breeding 5.0: Artificial intelligence (AI)-decoded germplasm for accelerated crop innovation. Journal of Integrative Plant Biology.
https://doi.org/10.1111/jipb.70008

7. Yu, H. (2026). AI foundation models in plant biology. New Phytologist.
https://doi.org/10.1111/nph.71395

8. Nasir, S., Khan, L., Murtaza, A., Raza, A., Nasir, F., & Azam, U. (2026). Bridging traditional plant breeding with precision agriculture: A path to smarter crop improvement. In Explicable Agriculture: Microbial Ecology, Bio-Innovations and Crop Production (pp. 194-201). Unique Scientific Publishers.
https://doi.org/10.47278/book.tl/2026.383

9. World Agriculture Forum. (2026). WAF Interactive Webinar on Biotechnology and AI Contributions to Plant Breeding [Video]. YouTube.
https://www.youtube.com/watch?v=xkVT9D7Nqy0

10. Prasanna, B. M. (2026). Predictive Breeding: Bridging Genetics, Artificial Intelligence and Agriculture [Video]. SHARE GENES.
https://www.youtube.com/watch?v=FWcxeEC90JA

11. Money Today. (May 16, 2025). More than 4,000 crop genomic data points and molecular markers: “K-Digital Breeding” gets underway.
https://www.mt.co.kr/economy/2025/05/16/2025051613191936623

4.2

Gene Editing (CRISPR) and Precision Breeding

In the greenhouse of the Department of Agricultural Biotechnology at the Rural Development Administration's National Institute of Agricultural Sciences, every pot bears a small label. The institution that conducted the 2024 OsNAC-edited rice study was the National Institute of Agricultural Sciences, not the National Institute of Crop Science. Reading the genes from a single cut leaf can reveal rice that is nearly identical to the original variety but has a different base at a predetermined position. CRISPR is a tool that selects a specific target and cuts or alters a gene in this way.

A plant's genes are long sequences of the four letters A, T, G, and C. Guide RNA directs the system to a target sequence of about twenty letters, and Cas9 cuts after checking the adjacent PAM sequence. It is easy to assume that the cut occurs at exactly one site, but that is not the case. Depending on the locations and number of mismatches and on the surrounding sequence, other partially similar sites may also be cut. These are called off-target effects. This is why guide design and whole-genome verification are necessary.

After cutting, the cell rejoins the broken DNA on its own. If it joins the ends without a template, a few letters may be deleted or added, disabling the gene. Homology-directed repair, which introduces a desired DNA template for precise correction, is also possible, but its efficiency is low in plants and varies widely by crop and tissue. Base editing changes a specific letter, while prime editing aims at a broader range of small corrections. Describing the process merely as filling in the exact letters when given a blueprint conceals both success rates and types of failure.

Editing tools enter cells through Agrobacterium, a gene gun, or RNP delivery into protoplasts. The RNP method, which combines a protein with guide RNA, has the advantage of producing edited plants without inserting foreign DNA. Even so, one cannot conclude before molecular verification that no foreign genetic material remains at all. Unexpected variations arising from the delivery vehicle and regeneration process must also be examined. Nanocarriers remain an experimental method being tested in several crops.

The problem is the enormous length of the rice sequence. The rice genome contains 400 million letters, and throughout it are regions that look almost identical to the twenty-letter target selected by the researcher. If the guide RNA binds to one of these similar regions instead of the true target and cuts at the wrong site, an unintended gene may be damaged or an unexpected trait may appear. It is impossible from the outset for a person to scan 400 million letters by eye and find every similar region. Artificial intelligence therefore takes on this work.

Wheat presents an even more difficult case. The wheat genome is hexaploid, with three sets of genomes from different ancestors layered together in one organism. The gene that causes susceptibility to powdery mildew is embedded in nearly the same form in all three sets, so crossbreeding offered no way to select and suppress just one set. The gene is named TaMLO, with copies called TaMLO-A1, TaMLO-B1, and TaMLO-D1 in the three genomes. Altering one copy still left the other two intact, and the plant remained susceptible to the disease.

CRISPR can send the same guide to each of the three sets and disable all three sites at once, and accurately finding and targeting those three sites simultaneously is the kind of calculation artificial intelligence performs well.

The history of gene editing and commercial crops must be distinguished by technology. The first commercialized high-oleic soybean in the United States was a product in which the FAD2 gene was edited with TALEN, not a CRISPR crop. The commercial Innate potato, engineered to resist browning, reduced PPO expression through RNA interference. Research continues on using CRISPR to alter disease resistance, leaf angle, and nutrient-use efficiency in maize and potato, but experimental CRISPR cases should not be grouped with TALEN and RNAi products already on the market as achievements of the same technology.

Artificial intelligence first scans the entire genome to find every region that bears even a slight resemblance to the target. A deep-learning model then assigns each candidate location a score indicating the probability that the molecular scissors will actually bind and cut there. These scores come from measured data, not guesswork. The model learns from tens of thousands of experimental results obtained by introducing the scissors into living cells and using fluorescent markers to identify where cuts actually occurred.

According to a paper by Abbaszadeh and Shalayi posted on arXiv in 2025, recent models using convolutional neural networks, recurrent neural networks, and transformer architectures identify sequences vulnerable to cutting far more accurately than earlier statistical formulas. These models provide more than a score: they also show which combinations of letters in a sequence increase the risk. Researchers can therefore examine and understand the reason with their own eyes.

The scissors themselves can also be changed. Using a modified enzyme with greater target specificity than standard scissors makes it less likely to bind to an unintended region in the first place. Researchers use this scoring table together with the new enzyme to place the lowest-risk combination on the laboratory bench, and they perform another round of verification after editing. They reread every gene in the plant and compare it letter by letter with the original genome, while artificial intelligence searches the comparison table for patterns too subtle for the human eye and checks for variations outside the target.

This whole-genome resequencing verification is costly even for a single line. Small breeding companies and research institutes in developing countries with limited budgets are tempted to skip this verification step, and some do proceed to the next stage without completing the final check before releasing plants into the field. Recent experiments have gone as far as using artificial intelligence models trained on the three-dimensional structures of proteins to design entirely new scissor proteins that do not exist in nature.

They are not yet ready for field use, but the role of artificial intelligence is moving from selecting the scissors to creating them.

Even after editing and off-target verification are complete, a plant cannot be released directly into a paddy. Researchers must determine whether the trait is inherited by offspring and then complete regional adaptation trials and variety registration. Speed-breeding greenhouses control light, temperature, and humidity to accelerate generations in crops such as wheat and barley. The responses of rice and maize vary by variety and photoperiod. Fixing a line quickly and completing a variety for sale to farmers require different lengths of time.

In 2024, Korea's Rural Development Administration announced that it had targeted some 2,600 rice defense-related genes and, in an initial phase, created 146 types of NAC-family edited plants. OsNAC30-edited rice showed a 51.9 percent reduction in bacterial blight lesions, while OsNAC59-edited rice showed a 24.5 percent lower mortality rate from bakanae disease. The researchers identified the two defense genes by directly infecting the edited population.

The regulatory gates that gene-edited crops must pass differ by country. In the United States, the 2020 SECURE rule was vacated by a court decision in December 2024, and APHIS reinstated its previous permit, notification, petition, and “Am I Regulated” procedures. The claim that the absence of foreign DNA allows a crop to proceed to field trials with a single sheet of paperwork is not a generally applicable rule in force in 2026. The European Union's Regulation (EU) 2026/1388 will apply fully from July 17, 2028.

The NGT1 criteria do not define equivalence to the original genome as no more than 20 differences; they consider the lengths of insertions and substitutions, the number of changes per coding sequence, and the total number of changes in the haploid genome together. Herbicide tolerance and traits that produce known insecticidal substances are excluded from NGT1.

Patents are another obstacle. The principal holders of foundational Cas9 patents are the UC Berkeley-Vienna group and the Broad-MIT group, while companies including Corteva have licensed some of the rights. Small breeding companies and public research institutes spend time and money determining which patents they need and negotiating licensing agreements. Securing the legal right to use the technology can take longer than learning the technology itself.

Somewhere between these regulatory systems, Japan has approved and now sells an edited tomato with higher gamma-aminobutyric acid content, claimed to lower blood pressure, as a non-genetically modified organism. Yet the fact that it is sold does not mean everyone buys and eats it. A 2025 survey by Philippine scholars found that among consumers who believed such products were safe, willingness to buy edited bananas increased by 267 percent in China and 171 percent in Japan. Conversely, reluctance toward unfamiliar foods reduced willingness to buy in all three countries. Three U.S.

universities reached a similar conclusion when they surveyed 1,638 people about low-acrylamide flour. People were considerably more comfortable with using

CRISPR on crops than on animals or humans. Concerns remained, however, that adverse effects might emerge later and that patents on the technology are held by a small number of large seed companies, making it difficult for small farms and researchers in poor countries to gain access.

Traits that pass all these procedures face one final gate. Rice that grows flawlessly in a greenhouse often fails to thrive in an actual paddy. In a greenhouse, people control temperature, water, and light; in a paddy, monsoons, typhoons, pests, and disease arrive together. A 2026 paper by the Iranian researcher Rahimi identifies why laboratory success fails to carry over to the paddy. Artificial intelligence models that select targets are usually trained only under normal conditions and have no prior exposure to stresses such as drought or extreme heat.

Even when the edit itself succeeds, researchers often proceed without examining how the gene interacts with other genes in the cell. Some publish after planting at only one or two sites, without repeated testing over multiple years in paddies across multiple regions.

A 2025 Cell paper inserted a 10-base heat-responsive element into the rice CWIN promoter. Under high-temperature conditions, the plot yield of the edited rice was 25 percent higher than that of Zhonghua 11, and it recovered up to 41 percent of the yield loss caused by high temperatures. The result of inserting a single short sequence was clear, but long-term adaptation trials repeated for several years across multiple regions remain to be done. More seasons must pass before anyone knows whether rice that performed well in greenhouse and plot trials will perform equally well in farmers' paddies.

Breeding companies and national research institutes are therefore extending the period in which they divide plants among multiple paddies and observe them for several years before announcing a greenhouse success. CRISPR can cut a single letter precisely, but whether that letter produces useful rice can be determined only through years of field trials. No one can yet guarantee how much longer the rice in that labeled pot will remain standing than the traditional variety in the neighboring paddy after the next typhoon passes.

Sources and References

1. Nasir, S., Khan, L., Murtaza, A., Raza, A., Nasir, F., & Azam, U. (2026). Bridging traditional plant breeding with precision agriculture: A path to smarter crop improvement. In Explicable Agriculture: Microbial Ecology, Bio-Innovations and Crop Production (pp. 194-201). Unique Scientific Publishers.
https://doi.org/10.47278/book.tl/2026.383

2. U.S. Department of Agriculture, APHIS. (2025). Vacatur of 2020 Biotechnology Regulations.
https://www.aphis.usda.gov/vacatur-2020-regulations

3. European Union. (2026). Regulation (EU) 2026/1388 on Plants Obtained by Certain New Genomic Techniques.
https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202601388

4. Fu, B. X. H., et al. (2022). Target-dependent Nickase Activities of the CRISPR-Cas Nucleases Cpf1 and Cas9. Nature.
https://www.nature.com/articles/s41586-022-04470-1

5. Lou, H., Li, S., Shi, Z., et al. (2025). Engineering source-sink relations by prime editing confers heat-stress resilience in tomato and rice. Cell, 188(2), 530-549.e20.
https://doi.org/10.1016/j.cell.2024.11.005

6. U.S. Food and Drug Administration. Genome Editing in Agricultural Biotechnology.
https://www.fda.gov/food/agricultural-biotechnology/genome-editing-agricultural-biotechnology

7. Rural Development Administration. (2024). Discovery of rice genes with high resistance to bacterial blight and bakanae disease.
https://www.rda.go.kr/board/board.do?dataNo=100000798693&mode=view&prgId=day_farmprmninfoEntr
y

8. ISAAA. (2025, July 30). Survey shows consumer acceptance of gene-edited bananas in the Philippines, Japan, and China. Crop Biotech Update.
https://www.isaaa.org/kc/cropbiotechupdate/ged/article/default.asp?ID=21404

9. ISAAA. (2025, February 26). Study reveals consumer acceptance of gene-edited food. Crop Biotech Update.
https://www.isaaa.org/kc/cropbiotechupdate/article/default.asp?ID=21220

10. ISAAA. (2021, September 22). Japan starts sale of genome-edited high-GABA tomato. Crop Biotech Update.
https://www.isaaa.org/kc/cropbiotechupdate/article/default.asp?ID=19024

11. Newsis. (October 2, 2024). Rural Development Administration discovers rice genes with high resistance to bacterial blight and bakanae disease. Newsis.
https://www.newsis.com/view/NISX20241002_0002907293

12. Sangwa, S., & Mutabazi, P. (2025). Artificial intelligence in agriculture: Ethical stewardship, responsible innovation, and governance for sustainable food systems. Intelligent Agriculture, 1(2), 25-45.
https://doi.org/10.54963/ia.v1i2.1684

4.3

Developing Climate Change-Resilient Varieties

Monsoon rain poured for three consecutive days onto a low-lying paddy not far from the Bay of Bengal. Young rice plants transplanted less than a month earlier are submerged beneath waist-high water. Stalks that stood green only three days ago are now invisible. Yet this paddy also contains rice that lifts its head again when the water recedes, rice whose genes have been altered to endure ten days underwater.

An international team of rice researchers examined 6,274 flood-tolerant germplasm accessions and selected 89 superior accessions, including 37 that showed strong tolerance. The study combined field selection, a submergence test lasting about three weeks, QTL analysis, and genomic selection. A one-way sequence in which machine learning first narrowed the candidates and tank tests then ran for 10 to 21 days differs from the original paper. The result that the 37 highly tolerant accessions had survival rates 40 to 50 percent higher than SUB1A introgression lines should be retained, but the selection process must be described correctly.

Whether near the Yellow Sea or the Bay of Bengal, paddies close to seawater face a different concern. Rice fails to grow because the soil itself is salty, not because water overflows it. On reclaimed land or in deltas affected by sea-level rise, white salt settles like frost on the dry paddy floor. As the roots draw up the saline water, they fail to filter for the nutrients they need, and the leaves burn inward from the tips. Like flood tolerance, salt tolerance cannot be resolved through a single gene.

The ability to distinguish among ions, withstand pressure in root cells, and conserve water is divided among dozens of different genes, and these genes must work together for rice to set panicles in saline soil.

Traditional breeding required these traits to be selected by eye. Breeders planted thousands of lines in saline paddies, waited several years, and then replanted only the survivors. During that time, every task of irrigating the fields, measuring salinity, and scoring the extent of leaf burn was performed by hand. This is why bringing a single variety into the world took eight to twelve years.

The idea of creating rice that thrives in saline soil is not new. In 2007, the Chinese agronomist Qifa Zhang (張 發) proposed in the international journal PNAS that genes for drought resistance and nitrogen-use efficiency scattered among wild relatives of rice be brought together in a single variety. He called it Green Super Rice. At the time, even tracking one gene with a marker through crosses took several years. Predictive breeding has now inherited that concept.

A computer first screens by calculation dozens of combinations of tolerance genes that had lain dormant in wild-rice gene banks, and only the combinations with high potential are tested in actual paddies.

When wheat encounters extreme heat during flowering and grain filling, pollen and grain formation are disrupted. A synthesis of studies from around the world estimated that wheat yields decline by about 6 percent on average for each 1°C increase in mean growing-season temperature. The 21 percent figure is an upper-end value reported under specific conditions and in particular studies, so it should not be presented as the response of all wheat fields. The vulnerable periods for winter wheat and spring wheat also vary by region and variety.

Breeders identify heat-tolerant lines by considering genetic markers together with temperature records from the flowering period.

Rice can experience sterility and loss of quality from high temperatures during the brief period when it flowers. Researchers at Japan's agricultural research organization used a weather grid of approximately 1 kilometer and a panicle-temperature model to compare combinations of early flowering, heat tolerance, and heading-date adjustment. They calculated how scenarios that advanced flowering by 3 hours or increased heat tolerance by 3°C would change the risk of sterility by region. The study appeared in 2026 as article 110519 in volume 344 of Field Crops Research.

Wheat and rice are not the only crops vulnerable to extreme heat and drought. When too little water rises through the xylem at the tips of soybean roots, the pods ripen without filling completely. Because soybean obtains nitrogen through symbiosis with root-nodule bacteria, drying soil around the roots breaks this relationship first. Leaves deprived of nitrogen lose color, and the beans shrivel. Predictive-breeding researchers scan the entire soybean genome to identify gene combinations that sustain root-nodule symbiosis and protein synthesis during drought, and they plant only lines containing those combinations in drought trial fields.

This brings the methods used for rice and wheat directly into soybean fields.

Flood-tolerant rice, salt-tolerant rice, and heat-tolerant wheat all face the same challenge: determining which combinations of genes will withstand the future climate takes too long. Traditional breeding had no choice but to plant crossed offspring in actual paddies and fields, expose them to drought, submerge them in water, or apply salt, and observe them for years. Artificial intelligence has transferred much of this waiting time into the computer. The method is called predictive breeding. It is also called genomic selection because the genome is read in advance to calculate the outcome.

DNA is extracted from the leaves of young seedlings, and tens of thousands of genetic markers are read at once. These markers indicate which genes an individual carries. Reading the markers makes it possible to estimate how well a seedling will withstand drought and heat when it grows, without raising it to maturity. The calculations employ several methods, including one that combines multiple decision trees and reaches an answer by majority vote, and artificial neural networks modeled on the neurons of the human brain.

Because markers and environmental data interact differently in each region, breeders layer several computational methods and select the one that best matches actual results. Records from a single year and location are insufficient. Calculations approach reality as records accumulate across years of severe drought and ordinary weather, and across sandy and clay soils. Adding weather and soil records accumulated over many years and locations makes it possible to select just a few hundred crosses worth planting in actual paddies and fields from tens of thousands of possible combinations.

Genetic markers are not the only data computers use in these calculations. Drones fly over trial fields, recording leaf color and leaf-surface temperature. The leaves of water-stressed rice are slightly warmer than those of healthy rice. The difference is invisible to the human eye but clearly captured by thermal cameras.

As these daily images and numbers are placed alongside genetic markers, the computer also learns how particular gene combinations appear in actual fields.

The DTMA project of the International Maize and Wheat Improvement Center released 60 drought-tolerant hybrids and 57 open-pollinated varieties between 2007 and 2012. The project targeted up to 40 million people in 13 African countries, not 40 million farming households. Under moderate drought, it aimed for a yield advantage of about 1 metric ton per hectare, or 20 to 30 percent. There is no evidence connecting this achievement to a reduction in the breeding cycle from ten years to three or four years through artificial intelligence prediction.

Speed breeding uses artificial lighting to control day length and temperature and accelerate generations. Results of 6 generations per year under 22-hour long-day conditions have been reported in crops such as spring wheat and barley. Many maize varieties respond to shortening days, so the same regimen cannot be applied to them without modification. Even if generation acceleration fixes a line quickly, multi-location cultivation trials and variety registration require separate time.

Accelerated breeding still has limitations. The data learned by computers reflect weather and soil conditions actually experienced over the past several years. Climate change, however, produces extreme heat and drought that no one has experienced before. The scores predicted by artificial intelligence become unstable when faced with extreme weather absent from the training data. If artificial intelligence rates tolerance higher than it really is under unprecedented heat, a farmer may trust that assessment, plant the variety, and lose the year's crop.

The procedure of planting computer-selected candidates in fields to determine whether they actually withstand the conditions has therefore not disappeared. Artificial intelligence reduces the candidates for planting from tens of thousands to a few hundred.

Computer calculations have barriers of their own. Processing tens of thousands of genetic markers together with years of weather and soil data requires high-performance computers and specialists capable of interpreting the results. Research institutes supported by international organizations, such as CIMMYT and IRRI, possess such equipment and personnel, but national breeding institutions in countries with limited budgets face different circumstances. If they cannot perform the calculations themselves, they cannot develop predictive models tailored to their own soil and climate and often borrow models developed in other countries.

These borrowed models have never been trained on the paddies and fields of the borrowing country, so their predictions retain the same instability when they encounter conditions absent from the training data. A new divide may open between countries with computational capacity and those without it.

The path from a testing station's refrigerator to a farmer's hands is long for a new variety. The informal share of seed used by African smallholders varies by country and crop. Data from Uganda estimate that informal channels, including neighbors, markets, and farmer-saved seed, account for approximately 85 to 89 percent. This cannot be fixed at more than 90 percent for all smallholders across Africa. Each region must be examined to distinguish cases in which farmers cannot plant because seed is unavailable, do not choose it because it is unfamiliar, or cannot buy it because it is too expensive.

Concentration in the U.S. seed market varies by crop. USDA estimates placed the shares held by the top two companies at approximately 71.6 percent for maize and 65.9 percent for soybean. The top four companies held about 94 percent of the cotton market. From 1990 to 2020, the seed price index, centered on GM seed, rose 463 percent, while the crop price index rose 56 percent. Seed prices rose much faster than the prices farmers received for selling grain.

A study by agricultural economists from several African countries, published in the international Journal of Agricultural Economics in 2026, demonstrates one way to lower this barrier. When new varieties were distributed in palm-sized small packs rather than large sacks, both understanding of the variety and the following year's purchase rate increased among recipient farmers. The International Seed Federation made a similar point in an article published in 2026.

Letting farmers watch a field first planted by a village elder, or having an agricultural extension officer visit each home and place a seed packet directly in the farmer's hand, works better than a single advertising leaflet. No matter how outstanding the trait, it will not spread unless farmers first trust it and plant it once.

We return to the low-lying paddy by the Bay of Bengal. Whether rice carrying the SUB1 gene actually lifts its head again after the water recedes can be confirmed only in that paddy, for that farmer, when the next monsoon arrives. The water still takes days to drain, and during that time it is the farmer, not the computer, who squats on the paddy bank and watches its color.

Sources and References

1. Prasanna, B. M. (2026). Predictive Breeding: Bridging Genetics, Artificial Intelligence and Agriculture [Video]. SHARE GENES.
https://www.youtube.com/watch?v=FWcxeEC90JA

2. Anumalla, M., Khanna, A., Catolos, M., Ramos, J., Sta Cruz, M. T., Venkateshwarlu, C., Konijerla, J., Pradhan, S. K., Dash, S. K., Das, Y., Chowdhury, D., Chetia, S. K., Das, J., Nath, P., Merugumala, G. R., Roy, B., Pradhan, N., Jana, M., Dana, I., ... Hussain, W. (2025). Future flooding tolerant rice germplasm: Resilience afforded beyond Sub1A gene. The Plant Genome, 18(2), e70040.
https://doi.org/10.1002/tpg2.70040

3. Zheng, Y., Cai, Z., Wang, Z., Maruza, T. M., & Zhang, G. (2025). The genetics and breeding of heat stress tolerance in wheat: Advances and prospects. Plants, 14(2), 148.
https://doi.org/10.3390/plants14020148

4. CIMMYT. (2026). Drought Tolerant Maize for Africa (DTMA). International Maize and Wheat Improvement Center.
https://www.cimmyt.org/projects/drought-tolerant-maize-for-africa-dtma/

5. BusinessWorld Online. (2024, March 27). Rice variety resistant to heat, drought to be launched in 2026, IRRI says.
https://www.bworldonline.com/economy/2024/03/27/584553/rice-variety-resistant-to-heat-drought-to-be-launched-in-2026-irri-says/

6. Mulungu, K., Subakanya, M., Setimela, P., Essegbemon, A., Chivasa, W., Gethi, J., Chiduwa, M., Sigalla, J., Mvungi, H. H., & Ngoma, H. (2026). Small seed packs, big potential? Effect of seed packs on knowledge and adoption of improved crop varieties. Journal of Agricultural Economics.
https://doi.org/10.1111/1477-9552.70040

7. International Seed Federation. (2026). Inclusive seed systems: Building trust with farmers. World Seed 2026.
https://worldseed.org/2026/06/03/inclusive-seed-systems-farmers-world-seed-2026/

8. Catholic Relief Services. (2026). Seed systems smallholder farmers use.
https://www.crs.org/sites/default/files/documents/2025-12/seed_systems_farmers_use_embargoed_until_jan_20.pdf

9. Zhang, Q. (2007). Strategies for developing Green Super Rice. Proceedings of the National Academy of Sciences, 104(42), 16402-16409.
https://doi.org/10.1073/pnas.0708013104

10. Toda, Y., et al. (2026). Simulation of Heat-Induced Spikelet Sterility in Rice and Evaluation of Adaptation Options Under Historical and Projected Meteorological Conditions in Japan. Field Crops Research, 344, 110519.
https://doi.org/10.1016/j.fcr.2026.110519

11. Dawson, J., Hubbard, K., & Jenney, P. (2026, June 23). How everyone pays the cost for patents on seeds, and private companies get rich from keeping them secret. The Conversation.
https://theconversation.com/how-everyone-pays-the-cost-for-patents-on-seeds-and-private-companies-get-rich-from-keeping-them-secret-284144

CHAPTER 5

AI Decision-Making and Automation

5.1

Agriculture-Specific Multimodal Large Language Models (Agri-LLM)

One afternoon, as the sun begins to sink, a yellowish spot is spreading across a rice leaf. The leaf was fine until yesterday. The farmer takes out his phone, photographs the spot, uploads the image to a consultation window, and asks the screen in a dialect-tinged voice, "What's causing this? Do I need to spray something?" If the answer is wrong, the stakes extend beyond a single leaf to the entire year's crop. The technology built to answer this question accurately is the agriculture-specific multimodal large language model, or Agri-LLM.

The farmer can upload the same photograph to an ordinary AI chat app on the phone and still receive an answer. The problem is that the answer is often wrong. General-purpose artificial intelligence is trained by memorizing vast amounts of human writing and stringing together plausible sentences. The difference between rice blast and bacterial leaf blight on a rice leaf may be a matter of only a few millimeters, even to the eye. A general-purpose model cannot distinguish that subtle difference.

Actually knowing the correct answer and composing a sentence that sounds correct are entirely different tasks, and general-purpose models excel only at the latter. If they do not know the name of a disease or pest or the proper treatment, they simply invent one. They present nonexistent pesticide names as genuine and confidently offer advice that does not fit the local climate or regulations. The AI industry calls this phenomenon hallucination. In farming, hallucinations cannot be laughed off. Spraying the wrong chemical at the wrong concentration can kill an entire field.

Agri-LLM is designed differently from the outset to reduce such hallucinations. It is retrained on a separate collection of specialized agricultural texts, including plant pathology papers, pest identification guides, and farm consultation records. A high-precision visual recognition system is added to extract, across multiple layers, even tiny lesions and the outlines of small insect larvae in photographs. It reads data from sensors that measure soil moisture and levels of nitrogen, phosphorus, and potassium, together with weather forecast signals.

It is called multimodal because it makes judgments by considering text, photographs, and sensor readings all at once. The pattern of spots extracted from a photograph does not immediately translate into the name of a disease. The model therefore retrieves a knowledge table in which disease names and their typical symptoms have been organized in advance, then compares the patterns in the photograph with them one by one. Only through this process of matching what it sees with what it already knows can a spot be turned into a disease name.

The moment the farmer takes a photograph, the phone gathers other information as well. It combines the satellite-determined location of the paddy, soil readings just transmitted by a nearby sensor, and a weather forecast for the next day. The system begins making a judgment only when the photograph, spoken words, and numerical data are all available. If one value is missing, it either postpones the judgment or provides only a provisional answer while identifying the missing item.

Before composing an answer, Agri-LLM first searches actual documents. This method is called retrieval-augmented generation (RAG). Think of a librarian. When asked a question, a librarian does not answer from memory alone. The librarian walks to the shelves, takes out an actual book, and points to the page and line where the information appears. Agri-LLM likewise first searches a database containing Rural Development Administration guidelines, plant pathology textbooks, and standard pest control guidelines.

Only after finding relevant documents does it formulate an answer based on their contents, and it displays a button linking to the source document on the screen. When the farmer presses that button, the original document can be checked instead of relying on the AI's words.

The performance of agricultural multimodal models is revealed through published test data and independent evaluations. A score for identifying diseased leaves, a score for detecting pests, and the quality of pest control advice come from different tests. Results change when the image background, crop variety, or evaluator's criteria change. The mere presence of numbers reported to decimal places does not mean that a model has been validated in paddies and fields. Before giving a farmer the name of a chemical, the data, comparison models, and failure cases must be disclosed together.

The picture becomes clearer when familiar names from rice farming are used. Rice blast, rice false smut, Fusarium head blight in wheat, northern corn leaf blight, and stem-boring leaf miner larvae all appear by name on the list of diseases and pests that Agri-LLM distinguishes, mirroring those that food-crop farmers encounter in practice. Research to test the capabilities of such models has followed. The AgMMU benchmark released by U.S. researchers in 2025 presents multiple AI models with crop photographs and questions simultaneously and compares the accuracy of their answers.

This means that the field has entered a stage in which multiple laboratories evaluate one another using the same yardstick, rather than relying on promotional material from a single company. The number of categories, 141, is not absolute in itself. It is meaningful, however, because it covers a broad range of diseases and pests that recur in the field rather than one or two laboratory cases.

AgriSmart, developed by researchers in Pakistan, combines crop recommendations, disease diagnosis, and a chatbot on a single screen. In the paper, crop recommendation accuracy reached 99 percent, while the disease model recorded an F1 score of 0.70 and an mAP@50 of 74 percent. These figures are model evaluation results obtained from a defined dataset. They are not field results showing that a farmer preserved the harvest in an actual paddy by delaying irrigation before rain and determining the amount of lime to apply. A full season of validation remains between a score on the consultation screen and the outcome in an actual field.

Kisan e-Mitra, operated by the Indian government with Wadhwani AI, provides information on government programs and agriculture in 11 regional languages. According to data from the Parliament of India, it had received 9.5 million queries from 5.3 million farmers by July 2025. Farmers who have difficulty reading can enter the consultation window by speaking into a microphone and listening to information about subsidies and insurance. These figures include people who spoke to the government by voice, rather than through a screen, for the first time.

Existing image-based disease and pest diagnosis services are in operation in South Korea. In a 2024 announcement by the Rural Development Administration, average accuracy exceeded 95 percent, human recognition accuracy was 95.3 percent, and the system covered 31 crops and 182 diseases and pests. The agricultural AI agent “AI Isak” is scheduled for service in 2027 under the official work plan. If the image diagnosis system currently in operation reads a single photograph, AI Isak aims to become the next point of access, connecting agricultural technology materials through conversation.

Yet some barriers remain beyond the reach of even the most sophisticated technology. One is language. Models trained on standard language often struggle with rural dialects and idioms. A sentence such as “What’s causing this?” is difficult for an interpreter who knows only a standard-language dictionary to understand when spoken in dialect. To overcome this barrier, India has added language-understanding models tailored to regional languages such as Hindi and Marathi, along with speech-recognition and speech-synthesis engines that work offline. This does not mean they understand every dialect.

They still sometimes respond with irrelevant questions when confronted with a strong accent or an uncommon idiom.

Text on a screen is another barrier for some people. In rural areas of developing countries and in rural South Korea alike, many older farmers cannot comfortably read text on a screen. A text chat window is useless to them from the outset. That is why Agri-LLM services consistently put voice first. Farmers ask by voice and listen to spoken answers. They do not have to stare at a screen, so a consultation can be completed by briefly straightening their backs and pausing their fieldwork. This is also why the services explain results as a person would speak, rather than displaying a screen filled only with numbers and tables.

Communication signals frequently drop out in mountain villages and remote fields. That is another barrier to overcome. A system that places a large model on a distant server and queries it every time cannot be used in such places. Developers therefore divide and compress models and install them directly on phones or small devices placed at the edge of a field. Recently received weather data and sensor readings are stored temporarily on the device so that it can continue making at least a minimal judgment even when the signal disappears completely. A small model running on the device knows less than a large model on a server.

When the signal is gone, however, limited knowledge is better than none at all.

Not every farm has welcomed the technology. Farmers have complained that diagnostic accuracy falls sharply when a smartphone camera shakes and produces a blurry photograph. Some stopped using the service after a few attempts because mobile data charges were burdensome. There have also been actual reports of artificial intelligence giving misguided advice because it failed to read the soil type or weather variables unique to a region correctly. It is still too early to say the technology is complete.

One question remains: who is responsible if the advice is wrong and ruins a year's crop? Until now, farmers have borne most losses caused by mistimed irrigation or delayed disease detection on their own. In February 2026, Dr. Sophia Duan of La Trobe University in Australia pointed out that too little discussion had addressed who bears the risk when artificial intelligence fails after becoming part of routine farm decision-making. She also said that unless developers and farmers have a structure for sharing risk, farmers will never trust the technology's accuracy, no matter how accurate it becomes.

This is also why Agri-LLM attaches a button linking to the source document for every answer. Evidence can be disclosed, but guaranteeing that it applies perfectly to this field in this season is another matter. The same is true for the farmer who asked the question in dialect. The consultation was free, but if a loss occurs, the farmer ultimately pays for it.

The same question applies to where information entered by farmers, such as field locations, yields, and transaction prices, ultimately goes. If this information accumulates on the servers of large technology companies and instead weakens farmers' bargaining power, consultation becomes more convenient while farmers' position worsens. One proposed way to reduce this problem is to keep information on the device or a local server and retrieve only what is strictly necessary. A 2025 study by Sangwa and Mutabazi proposes design principles to reduce these risks.

It argues that value-sensitive design reflecting farmers' circumstances should be incorporated from the system design stage and that agricultural experts should regularly score the quality of answers to filter out errors missed by others. In an earlier warning issued in 2022, the University of Cambridge's Centre for the Study of Existential Risk noted that artificial intelligence pursuing yield alone could damage soil and streams in the long run. A standard is needed that measures the condition of the land several years later as well as the immediate yield. No structure yet exists in which the technology assumes responsibility on its own.

This development is poised to go one step further. If Agri-LLM is connected to a digital twin that reproduces an entire farm inside a computer, a farmer will be able to ask in advance how much yield would decline if a drought continued for the next 2 weeks. The model would run various weather scenarios in a virtual field and translate the results into language people can understand. Once it reaches the stage of issuing direct instructions to drones and automatic paddy sluice valves, the model that offered advice will physically operate the farm.

The moment a single system takes charge of both judgment and execution, the question of responsibility becomes far more serious. This stage, however, remains at the level of research and demonstrations. Few farms entrust such autonomous decisions directly to actual paddies and fields.

We return to the paddy bank. Morning sunlight spreads across the rice leaves, and the farmer reads the answer on the screen. A disease name and a chemical name appear, with a small button below them linking to the source document. The farmer's finger pauses over the button. Will the farmer press it before spraying?

Sources and References

1. Parliament of India. (2025). Kisan e-Mitra Usage and Language Coverage, Unstarred Question No. 1464. https://sansad.in/getFile/annex/268/AU1464_CSuYc2.pdf?source=pqars

2. Arun, T., Pooja Sri, M., & Priya, K. (2026). FarmGPT: An explainable AI powered smart agriculture system for crop recommendation and decision support. AIJFR, 7(3). https://doi.org/10.63363/aijfr.2026.v07i03.6414

3. Baloch, S., Ali, A., & Naeem, A. (2025). AgriSmart: An AI-powered farming assistant and crop management platform for precision agriculture. International Journal of Innovative Science and Technology, 7(10), 155-162. https://doi.org/10.33411/ijist/1710

4. Farmers Weekly. (2026). On-farm AI: New tech must share the risk. https://www.farmersweekly.co.nz/technology/on-farm-ai-new-tech-must-share-the-risk/

5. Gauba, A., Pi, I., Man, Y., Pang, Z., Adve, V. S., & Wang, Y.-X. (2025). AgMMU: A comprehensive agricultural multimodal understanding benchmark. arXiv.
https://arxiv.org/abs/2504.10568

6. Sangwa, S., & Mutabazi, P. (2025). Artificial intelligence in agriculture: Ethical stewardship, responsible innovation, and governance for sustainable food systems. Intelligent Agriculture, 1(2), 25-45.
https://doi.org/10.54963/ia.v1i2.1684

7. Rural Development Administration. (2026). 2026 Major Work Plan: Advancement of the Agricultural AI Agent AI Isak, with Service Scheduled for 2027.
https://www.rda.go.kr/board/board.do?mode=html&prgId=pln_mainbizQuery

8. University of Cambridge. (2022). Risks of using AI to grow our food are substantial and must not be ignored, warn researchers.
https://www.cam.ac.uk/research/news/risks-of-using-ai-to-grow-our-food-are-substantial-and-must-not-be-ignored-warn-researchers

9. Rural Development Administration. (2025). Increasing Farm Income by 20% with Artificial Intelligence (AI): Strategy for Converging Agricultural Science and Technology with Artificial Intelligence (AI). Rural Development Administration Press Release.
https://www.rda.go.kr/board/board.do?mode=view&prgId=day_farmprmninfoEntry&dataNo=100000806197

10. YTN. (2024). Diagnose and Treat Crop Disease and Pest Damage with a Smartphone. YTN Science.
https://m.science.ytn.co.kr/program/view_today.php?s_mcd=0082&key=202409261626341200

5.2 Agricultural Robotics and Task Automation

Morning fog hangs low over the Mangyeong Plain in Gimje as a tractor starts its engine. No one is in the driver's seat. The steering wheel moves slightly from side to side on its own, and an umbrella-shaped antenna on the roof blinks as it receives satellite signals. A farmer standing at the edge of the field looks only at the phone screen in hand, making no move to take the wheel.

There is no single reason this tractor is moving on its own. Across the paddies and fields where rice, wheat, maize, and soybeans are grown, fewer people are available to work. The village's young people who once transplanted rice together have left for the cities, and those who remain have reached an age when their backs are bent and their knees ache. It is now difficult to find anyone willing, even for pay, to bend beneath the blazing sun from dawn to dusk, transplanting seedlings, spraying chemicals, and carrying grain.

Labor costs rise each year, while infrastructure such as village halls and agricultural machinery repair shops closes down one facility at a time. A tractor plowing a paddy on its own is a choice made to fill the place left by people who are no longer there.

All these processes repeat at a speed the human eye cannot follow. Four stages run continuously in milliseconds: cameras and sensors perceive what they have just seen, the machine recalculates its position, redraws its route, and issues commands to its wheels and arms. In the past, a person drove farm machinery by holding the steering wheel and watching the surroundings. The machine now performs all four stages by itself.

Two systems keep this tractor moving straight along the paddy bank. One is a satellite in the sky; the other is a camera mounted on the front of the tractor. When only global navigation satellite system (GNSS) signals are used, positional error can approach 1 meter. The tractor therefore also uses real-time kinematic satellite navigation (RTK GNSS), which receives correction signals from a ground reference station in real time. When South Korean agricultural machinery manufacturer Daedong put one of its tractors through a national autonomous-driving test in 2023, its lateral deviation from the planned route remained within 7 centimeters.

The company said it would release a Level 4 unmanned tractor with no person aboard by 2026. Spanish startup Voltrac said it would commercially launch a fully electric unmanned tractor in 2026, while South Korean agricultural technology company GINT continues field demonstrations of autonomous tractors and rice transplanters in Vietnamese paddies.

Autonomous agricultural machinery is equipped with a satellite antenna, an inertial measurement unit, cameras, and obstacle sensors. In 2026, one four-wheel vehicle drove a 30-meter section of a maize field 10 times and recorded an average lateral error of 6.2 centimeters and driving accuracy of 95 percent on cloudy days. The test focused on its ability to follow crop rows. Variable-rate seeding, which changes planting depth and spacing by zone, requires a separate control device and germination tests.

The wheels are no ordinary wheels either. With four-wheel drive and four-wheel steering (4WD/4WS), in which each wheel turns and steers independently, the vehicle can avoid slipping and turn in place even in the sticky mud of a flooded paddy or on a sloping field. Even in the middle of a flooded paddy where a person's feet would sink with every step, the robot maintains its balance by distributing a different amount of power to each of its four wheels.

If the satellite signal briefly disappears beneath a bridge or in the shadow of a windbreak, the camera takes over. It follows the rows of newly emerged rice seedlings in the nursery bed and identifies the centerline between them. The programs that perform this work have their own names. In rice paddies, a program called Fast-SCNN divides an image into small sections and classifies each section as seedlings or water, as though coloring them separately. In maize fields, another program called ST-YOLOv8s relearns leaf shapes as they change at each stage of growth so that it does not lose the centerline of a row.

Both programs redraw an invisible line in the image at every moment.

An autonomous rice transplanter that combines the two sensing systems kept row deviation within 6.2 centimeters at the demonstration site. In the field, muddy water splashes onto the lens and the midday sun shines directly into it. If light from LiDAR reflects off wet rice leaves and bounces away, the robot can momentarily lose track of its position. Higher water- and dust-resistance ratings are therefore required, together with calibration that realigns the readings whenever the camera and satellite values diverge.

An unmanned weeding robot demonstrated jointly by Japan's National Agriculture and Food Research Organization (NARO) and Newgreen runs on solar power and stirs the bottom of a flooded paddy with chains, clouding the water. If sunlight cannot penetrate the water, weeds such as barnyard grass cannot germinate. In an organic paddy, the dry weight of weeds fell from 35.4 grams per square meter to 5.6 grams. Rice yield in the same paddy rose by 20 percent.

Designs for weeding robots that combine an inexpensive camera, a Raspberry Pi, and MobileNetV2 have also been proposed. One design published in IJRASET combined four wheels, a camera, a robotic arm, and cloud-based records in a single unit, but it had not reached the stage of building an actual robot and testing it in a field. Prototypes and field trials must determine whether this design can withstand dust, mud, and moving leaves.

Crop-spraying drones can be programmed to plan routes that avoid covering the same area twice. The SprayCraft study tested a path-planning algorithm using synthetic location data. It found a short route on a computer screen, but this was not a flight that withstood wind, battery constraints, and the weight of a chemical tank in a real field. Equipment that reads disease density with a camera and adjusts nozzle output must undergo testing separate from that of the path-planning algorithm.

Heads filled with mature grain are more difficult to handle. Whether the head is maize or rice, slightly too much gripping force can burst the kernels or damage their skins, while too little force allows them to slip and fall. Robotic arms are equipped with both pressure sensors and slip-detection sensors to maintain this balance. In an outdoor orchard experiment with a robotic arm that picked apples, the damage rate was 20 percent when the arm gripped fruit with its force-feedback system turned off. With the same system turned on, the damage rate fell to 0 percent. The robot itself must sense how hard its fingertips are pressing.

Rice and wheat are generally cut whole and fed into a thresher, but this sense of touch matters in tasks that require selecting individual heads for seed or removing individual grains.

Robots that transport harvested crops follow workers across sloping fields or carry boxes to designated locations. A case reported by Japan's Ministry of Agriculture, Forestry and Fisheries found that labor fell by about 50 percent and physical strain by about 30 percent. The source does not report a reduction in labor intensity of more than 70 percent. The time people spend lifting loads and the time they spend supervising robots must be compared using the same standard.

The data generated by all these robots are useful only if they are sent somewhere. The robot's location, battery level, and most recently completed task are uploaded to a real-time web interface using Firebase, and if a problem occurs, a messaging bot, Telegram Bot, displays a short sentence on the farmer's phone. This interface is built by connecting a Python-based web application, Flask, to a real-time database. If something goes wrong, a person can press a single button on the screen to tell the robot to redo that section. Even when the robot appears to act entirely on its own, a person is still watching from behind the screen.

Autonomous agricultural machinery needs systems that detect people, animals, and obstacles and stop the machine if it leaves a designated area. ISO 18497-1 and 18497-2 address principles for machine design and obstacle protection. U.S. Sugar deployed 5 autonomous tractors across a 255,000-acre operation and ran them 24 hours a day. The 5 tractors did not handle the entire expanse of land. People left the driver's seat and moved to a control station where they monitored several screens at once.

Human hands are still needed where paddies and fields have irregular shapes. On terraced paddies where drainage channels turn sharply or paddy banks wind and curve, satellite signals and camera recognition often conflict, causing robots to lose direction and stop. In such places, a person still climbs onto the tractor and takes the wheel. Automation first gained a foothold in straight, regularly shaped paddies and fields.

Results collected from farms that introduced these robots show some recurring figures. Agricultural labor for tasks that once required constant human effort fell by 70 to 80 percent. Reports followed that inputs poured onto paddies and fields, such as fertilizer and pesticides, fell by 30 to 50 percent. These figures were pronounced in tasks such as weeding, pest control, and transport, which workers had endured with their backs and waists. The problem comes next.

RaaS, which allows farmers to subscribe to robots rather than buy them, is one way to spread the initial cost. A market report projects that the all-industry RaaS market will grow from $32.08 billion in 2026 to $67.85 billion in 2030. This total includes logistics, cleaning, and security robots, so it cannot be treated directly as the size of the agricultural robotics market. In rural areas, a more realistic approach is for village rental centers and cooperatives to establish an order for planting and harvest seasons and share a single machine.

The results to date reveal a clear boundary. Robots outperform people at identical, repetitive work in large, regularly shaped paddies and fields: following rows, applying a set amount of material, and moving a set weight. They do not get hurt or tired, and they do not rest at night. Work that requires distinguishing ripe grain from unripe grain by sight or making an immediate judgment when confronted with an unfamiliar obstacle still belongs to people.

Researchers are already envisioning the next stage. A system in which robots understand human speech, make their own judgments, and move their bodies is called embodied AI. Under this vision, several robots divide roles and act on the basis of a single-sentence instruction to pull weeds in Zone A and move harvested crops from Zone B. Researchers are also studying systems in which a crop-spraying drone first surveys a field and draws a disease and pest map, after which a ground-based weeding robot and spraying rover travel directly to those coordinates.

They even envision cameras and drones patrolling fields in the middle of the night while people sleep and compiling a list of tasks for the next morning. These systems remain at the paper and prototype stage. Whether this vision will work on actual paddy banks will become clear only after it fails several more times amid dust, mud, and torrential rain.

Once the sun is fully up, the unmanned tractor moves to the next paddy. The farmer still stands on the paddy bank, checking the remaining battery charge and unfinished section on the phone screen. While the tractor tills the soil and aligns the rows, the farmer uses both hands to open the sluice in the neighboring paddy.

Sources and References

1. Alda , M. C., & Eker, B. (2026). Development and field evaluation of an autonomous four-wheel-drive agricultural vehicle tracking crop rows using computer vision technology. Journal of Agricultural Production, 7(1), 13-19.
https://doi.org/10.56430/japro.1826810

2. Asami, H., & Nakamura, T. (2026). Effects of an Automatic Weed Suppression Robot for Paddy Fields in a Semi-Mountainous Area on Weed Suppression and the Growth and Yield of Paddy Rice. Japanese Journal of Crop Science, 95(1), 28-35.
https://doi.org/10.1626/jcs.95.28

3. Chen, K., Li, T., Yan, T., Xie, F., Feng, Q., Zhu, Q., & Zhao, C. (2022). A soft gripper design for apple harvesting with force feedback and fruit slip detection. Agriculture, 12(11), 1802.
https://doi.org/10.3390/agriculture12111802

4. International Organization for Standardization. (2024). ISO 18497-1:2024 Agricultural machinery and tractors, safety of partially automated, semi-autonomous and autonomous machinery, part 1: machine design principles and vocabulary. ISO.
https://www.iso.org/standard/82684.html

5. International Organization for Standardization. (2024). ISO 18497-2:2024 Agricultural machinery and tractors, safety of partially automated, semi-autonomous and autonomous machinery, part 2: design principles for obstacle protection systems. ISO.
https://www.iso.org/standard/82686.html

6. Ayalew, H., et al. (2024). SprayCraft: Efficient Coverage Path Planning for Agricultural Spraying. arXiv.
https://arxiv.org/abs/2412.12176

7. Alatan Tana. (2026). Analysis of the Impact of Intelligent Agricultural Machinery Retrofitting on Agricultural Production Efficiency. Agricultural Science, 9(2), 114-116.
https://doi.org/10.32629/as.v9i2.3755

8. Daedong. (2023). Daedong’s Autonomous Agricultural Machinery Becomes the First in the Industry to Pass the National Autonomous-Driving Test. Daedong News.
https://ko.daedong.co.kr/news/daedongnews/1485

9. Precision Farming Dealer. (2026). U.S. Sugar launches largest deployment of autonomous tractors in sugar industry.
https://www.precisionfarmingdealer.com/articles/7129-us-sugar-launches-largest-deployment-of-autonomous-tractors-in-sugar-industry

10. Future Farming. (2026). Voltrac's fully electric autonomous tractor to launch in 2026.
https://www.futurefarming.com/tech-in-focus/autonomous-semi-autosteering-systems/voltracs-fully-electric-autonomous-tractor-to-launch-in-2026/

11. Research and Markets. (2026). Robotics as a Service (RaaS) market report 2026.
https://www.researchandmarkets.com/reports/5896010/robotics-service-raas-market-report

12. Seoul Economic Daily. (2026). GINT Demonstrates Autonomous Operation of Tractors, Rice Transplanters, and Other Agricultural Machinery in Vietnam. Seoul Economic Daily.
https://www.sedaily.com/NewsView/2K7EFUEF1V

5.3

AI Tools Tailored to Smallholder Farms

One majigi of paddy is approximately 200 pyeong. Early in the morning, a person crouching on the bank of that flooded paddy holds a smartphone with a cracked corner on its screen, while the signal indicator shifts between one bar and two. A brown spot began spreading across a rice leaf overnight, and the person points the camera at the dew-covered leaf. There is no drone, no satellite imagery, and no soil sensor. All the person has is the phone in a pocket and a single app installed on it.

Roughly 80 percent of farms in developing countries are smallholdings of this scale. Yet while U.S. maize fields yield more than 10 metric tons per hectare, yields in sub-Saharan Africa remain between 2 and 3 metric tons. The term precision agriculture was originally coined with large farms equipped with autonomous tractors costing hundreds of millions of won or drone-mounted hyperspectral cameras in mind. Such prices have never entered the calculations of a person farming a single majigi of paddy. For someone farming at this scale, a useful tool is one that runs on the phone already in hand.

Plantix, developed in Germany, is a free app that identifies diseases from uploaded photographs of leaves, and it ranked first worldwide in downloads among agricultural apps. In Kenya, Farmer.Chat, developed by Digital Green, opened as a service through which farmers could ask and answer questions about diseases and pests as if speaking in a chat window. In its first year alone, it received more than 300,000 questions from 15,000 users. In Bangladesh, KrishokBondhu emerged as a voice consultation service that responds in Bengali when addressed in Bengali, allowing farmers who cannot read to seek advice as easily as making a telephone call.

All three are free of charge, and all three were built with one-majigi farms as their first customers. They are designed to accept any method a farmer can use, whether photographs, text, or voice. Several similar tools, including SCA-MobiPlant and PlantShield AI, have followed.

Kaohsiung City's "農來訊" in Taiwan is a public service that provides weather, disaster, agricultural commodity price, and IoT information. Rather than automatically controlling paddy sluice gates, it serves as a portal that gathers scattered information for farmers on a single screen. PlantVillage Nuru is a separate mobile tool that assists in diagnosing potato and sweet potato diseases. The EfficientPNet study tested its model on PlantVillage potato images with clean backgrounds. A public alert app, an in-field disease diagnosis system, and a laboratory image model may all fit on the same phone, but they have different performance records.

What comes after receiving the answer on the screen can be an even greater obstacle. To buy the prescribed chemical, farmers often have to ride a motorcycle to an agricultural chemical dealer in a town several stops away, and they must cut other spending to afford it. The diagnosis was free, but farmers must still bear the travel and expense required to follow the prescription.

Running a model on a low-cost Android phone requires reducing the amount of computation. Quantizing FP32 values to INT8 usually reduces parameter storage to about one-quarter. INT8 conversion alone does not shrink a model to a few hundredths of its size. TensorFlow Lite is a mobile runtime framework, while ONNX is a model exchange format. Vosk is a speech-recognition tool, and mBERT is a text encoder. These different tools must not all be grouped together as compression techniques or speech-recognition models. After compression, the model must be tested again to determine how much accuracy declines for rare diseases and regional accents.

PxD's evaluation of Ama Krushi in the Indian state of Odisha found that average yields among farms receiving the service rose by 4.1 percent. The benefit-cost ratio of 6 to 18 is a range calculated from program costs and estimated benefits, not net profit directly verified in every farm's bank account. A 2025 meta-analysis in Global Food Security reported that digital agricultural advisory services raised yields and income by an average of 6 percent each, but many individual estimates were not statistically significant. Average effects must be read separately from confirmed returns for individual farms.

When KrishokBondhu was tested with questions from actual farmers, it produced answers that human listeners considered satisfactory for 72.7 percent of a range of farming questions. When it was scored alongside the previous text-only method, KrishokBondhu received 4.53 out of 5, surpassing the other method's score of 3.13. For farmers with limited reading and writing skills, this difference matters more than any other feature. They can simply speak instead of tapping the screen to search through menus.

According to 2026 GSMA data, mobile internet use in low- and middle-income countries was 64 percent among women and 73 percent among men. The absolute difference is 9 percentage points, while 13 percent is the relative gap calculated against the male usage rate. Smartphone ownership and mobile internet use are different indicators as well. Field research continues to show that women farmers have less access to phones, mobile data, and time for training, but it must be made clear which indicator each gap represents.

Older farmers face another barrier. The ability to ask by voice and receive a spoken answer lowers the threshold for them as well. Yet people who have read paddies with their eyes and hands for more than 40 years often do not immediately trust an answer produced by a machine. Some services display the source of the guideline supporting a prescription beside it. Identifying a national agricultural research institute guideline or a standard crop protection product document allows farmers to check the basis immediately. Showing one more source is still not enough.

Actual usage rates differed markedly between villages where a field extension worker stood beside farmers and showed them how to press the controls several times and villages where no one did.

Efforts are underway to close this gap. Government agricultural technology centers or extension workers from local agricultural cooperatives may gather people at a village hall and hold a session where they practice together, beginning with how to hold the phone and press its controls. Such field education takes far more effort than updating a screen, but its effects last longer. Multiple field reports state that actual adoption rates differed between villages where a company left behind only an app and withdrew and villages where a local organization remained involved for several months.

Reports from multiple sites also repeatedly indicate that villagers trying the service first and telling one another by word of mouth is more powerful than any advertisement.

There is no guarantee that a service launched for free will remain free. Plantix is one example. It began as an app that diagnosed diseases to reduce pesticide use, but circumstances changed when investors came in. It expanded into a business that bought pesticides and fertilizers in bulk from companies such as Dow, BASF, Syngenta, and Bayer and resold them to the very retailers farmers visited. In 2023, the company itself was acquired by the chemical manufacturer Helm. The screen where farmers upload leaf photographs is still free.

Behind that free service, however, there is now a separate calculation that determines which pesticide to recommend first.

Another concern is that paddy locations, soil readings, and yield records entered by farms accumulate unchanged on corporate servers. A company holding that information could later use it as grounds to lower the price when buying those agricultural products. Alternatives under discussion therefore include a local government directly funding and taking responsibility for an app, as in Kaohsiung, or a cooperative such as an agricultural cooperative sharing operational responsibility while explicitly establishing that the data belong to the farm.

Instead of charging a price, such a structure covers operating costs through taxes or cooperative dues.

Even among free services, those sustained by taxes and donations and those under pressure to repay investors move in different directions over time. Just as large farms rent robots or drones through Robotics as a Service (RaaS), an advisory service for smallholders can endure only if someone ultimately pays the fee on their behalf. Farmer.Chat and KrishokBondhu are still supported by foundations and research funding. No one can guarantee how today's free service will change if that money ends.

The next stage under discussion is to install an ultralight language model directly on the phone, allowing natural conversational consultation to continue even in remote areas with no internet connection at all. For now, this remains closer to a laboratory demonstration, and it will take time before it reaches the hands of someone farming a single majigi of paddy. In the meantime, the signal will continue to shift between one and two bars, and the phone will still cost 100,000 won.

We return to that paddy bank. A few seconds after uploading the photograph, the person reads the answer that appears on the screen. A small advertisement appears beside the sentence explaining when and how much chemical to spray. The person cannot recall exactly when the advertisement first appeared there. Still, the person will trust this answer again and adjust the sluice. Whether the screen will remain the same next season depends on the circumstances of the company operating the service.

Sources and References

1. Majiwa, E., Ndaka, A. K., Ratemo, H., & Agembo, E. (2026). Artificial intelligence (AI) adoption and its perceptions among smallholder women farmers in Kenya. Discover Agriculture, 4(1), 162.
https://doi.org/10.1007/s44279-026-00626-z

2. Tesfay, W. (2026). Impact of improved irrigation technology adoption under groundwater irrigation and its impact on smallholder farmers' crop productivity, income, and food security in Southern Tigray Zone, Ethiopia. Research Square (preprint).
https://doi.org/10.21203/rs.3.rs-9797920/v1

3. Beach, R. H., Milliken, C., Franzen, K., & Lapidus, D. (2025). Meta-analysis of the impacts of digital information interventions on agricultural development. Global Food Security, 45, 100866.
https://doi.org/10.1016/j.gfs.2025.100866

4. Mutuku, B., Akuffobea-Essilfie, M., & Kansiime, M. (2026). Understanding gender disparities in access to digital advisory services among women and men farmers in Ghana. Frontiers in Sustainable Food Systems, 10, 1614245.
https://doi.org/10.3389/fsufs.2026.1614245

5. Omotayo, A. O., & Omotoso, A. B. (2026, June 3). AI offers promise for agriculture, but smallholder farmers risk being left behind. The Conversation.
https://theconversation.com/ai-offers-promise-for-agriculture-but-smallholder-farmers-risk-being-left-behind-279013

6. Singh, N., Wang'ombe, J., Okanga, N., Zelenska, T., Repishti, J., G K, J., Mishra, S., Manokaran, R., Singh, V., Rafiq, M. I., & Gandhi, R. (2024). Farmer.Chat: Scaling AI-powered agricultural services for smallholder farmers. arXiv.
https://doi.org/10.48550/arXiv.2409.08916

7. Ameen, M. R., Islam, A., Aktar, F., & Rafat, M. S. (2025). KrishokBondhu: A retrieval-augmented voice-based agricultural advisory call center for Bengali farmers. arXiv.
https://arxiv.org/abs/2510.18355

8. Miller, S. R. (2024, October). This app set out to fight pesticides. Once VC stepped in, the app helped sell them. Food and Environment Reporting Network.
https://thefern.org/2024/10/this-app-set-out-to-fight-pesticides-once-vc-stepped-in-the-app-helped-sell-them/

9. Precision Development (PxD). (2025, February 26). Customized digital advice can help farmers reduce crop loss and manage weather shocks: A summary.
https://precisiondev.org/customized-digital-advice-can-help-farmers-reduce-crop-loss-and-manage-weather-shocks-a-summary-or-as-much-as-we-can-summarize/

10. GSMA. (2026). The mobile gender gap report 2026. GSMA.
https://www.gsma.com/gender-gap/

CHAPTER 6

Climate-Smart Agriculture and Sustainable Ecosystems

6.1

Climate-Smart Agriculture (CSA) and Soil Ecosystem Management

Take a handful of paddy soil dug up with the tip of a hoe. Invisible organisms live in the black earth crumbling in your palm. Bacteria, fungi, and threadworms intertwine to break dead rice straw into small pieces, produce nitrogen for roots to absorb, and band together to block the way when unfamiliar pathogens enter. Countless organisms live together in a single handful of soil.

This soil has gradually deteriorated over the past several decades. Each time a tractor plowed it, the air pockets between soil particles collapsed, and each application of chemical fertilizer reduced the work done by microorganisms. As a result, the soil lost its ability to retain carbon, while paddies and fields increasingly became sources of carbon dioxide, methane, and nitrous oxide. The farming practices intended to reverse this trend are called Climate-Smart Agriculture. The name may sound grand, but the goals are clear: maintain yields, withstand drought and extreme heat, and restore carbon to the soil at the same time.

The international community describes these goals as three pillars: increasing productivity, strengthening resilience to climate shocks, and reducing greenhouse gases. Accomplishing all three at once is bound to be difficult.

When a paddy is filled with water, oxygen disappears from the soil. Microorganisms do not stop working even where there is no oxygen. They decompose rice straw and organic matter in a different way, producing methane as a by-product. Methane has a far greater capacity to warm the atmosphere than carbon dioxide, so the world's rice paddies have long been identified as a major source of greenhouse gas emissions. As evidence has shown that this problem can be addressed through water, water management technology is becoming standard practice in rice cultivation.

The practice of deliberately draining a paddy to let the soil dry briefly before reflooding it, instead of keeping it continuously flooded, is called midseason drainage.

According to the Rural Development Administration, multiple drainage, in which water is removed from paddies several times during the growing season, reduced methane by up to 44 percent compared with continuous flooding, while dry-field harrowing before rice transplantation produced a 14 percent reduction. The figures indicating an expansion from eight sites last year to 12 sites and more than 60 hectares this year refer to the plan to expand the dry-field harrowing program, not multiple water management and ICT as a whole. Reduction rates must be matched to the technology being deployed.

Sensor records can be used to verify compliance with water management practices, but it should not be asserted that the same equipment was installed at every demonstration site.

Not everyone immediately welcomed midseason drainage. For farmers who had long believed that paddies had to remain fully flooded to suppress weeds and ensure healthy rice growth, intentionally draining the water felt like a dangerous gamble. Concerns remain that rice plants might wilt during the few days when the paddy floor is dry or that weeds might seize the opportunity to emerge. To open and close the sluice based solely on the water level reported by a gauge, farmers must temporarily set aside the intuition ingrained in their bodies over many years. Crossing that threshold takes more than numbers and requires time.

The changes brought by midseason drainage do not end with methane. Paddies where the practice was followed used 20 to 30 percent less water, reducing by the same amount the water that had to be drawn from groundwater and rivers. In an experiment that inserted a fallow year into paddies otherwise planted with rice every year and used artificial intelligence to track the movement of nitrogen in the soil, the efficiency with which rice absorbed organic nitrogen from the soil increased noticeably. The rice absorbed more of the fertilizer applied, and less unabsorbed nitrogen was washed away by rainwater to pollute rivers.

The soil is compacted wherever a tractor passes through a paddy. Each pass of its heavy wheels compresses and eliminates the spaces between soil particles, and roots unable to penetrate the hardened layer grow sideways around it. There is no way to tell by sight where or how severely the soil has been compacted. A method has therefore begun to be used in which an electromagnetic induction sensor is pulled across the ground to send an electric current through it and measure the soil's electrical conductivity, after which the readings are fed into a machine-learning model to inversely calculate the degree of compaction at each depth.

Just as a mine detector locates metal underground, this sensor finds the location of invisible hard layers. If an implement that breaks up the subsoil is used only in the severely compacted sections identified on the resulting map, paths can be opened for roots without plowing the entire paddy. This method is not perfect, however. The values read by an electromagnetic induction sensor are affected by soil moisture and salinity, so a measurement taken the day after rain may differ from one taken during dry weather. The soil must therefore be measured again as the seasons change rather than only once.

Fields planted with wheat and corn call for a different prescription. This approach combines no-till farming, which minimizes plowing and leaves harvested stalks and leaves on the soil without turning them under, with charcoal-like biochar and a silicon-based soil amendment. Biochar settles between soil particles and retains water. In the U.S. Corn Belt, this amendment technology was demonstrated by overlaying subsurface moisture-sensor data with satellite time-series data to identify yield gaps that varied by zone and concentrating resources in the sections with the largest gaps.

Even in a year of severe drought, yields fell less in fields treated with biochar and silicon.

A different approach is used on steep mountain fields in places such as Guizhou Province in southwestern China. A fertigation system combines maps showing variations in terrain elevation with soil-moisture sensors to deliver water and fertilizer precisely and only near the roots, preventing soil from being washed away by monsoon rains and reducing the loss of nutrients flowing into waterways below the slopes. Yet producing biochar and transporting it to fields is costly, and its quality varies depending on the types of wood and rice straw used as feedstock, so in some years it does not produce the expected results.

The traditional dry-combustion method for measuring soil carbon burns soil at high temperatures and measures the carbon dioxide released. The method that estimates organic matter from the difference in the weight of the ash remaining after combustion is the loss-on-ignition method. The two methods must not be described as the same. Even within one field, values differ between a corner treated with manure and a gravelly slope, so samples must be collected from multiple locations, while the number and depth of samples are determined according to statistical design and validation criteria.

Sensors and satellites help select soil-carbon survey locations and estimate changes across broad areas. Cases introduced by PatSnap report an R^2 of 0.83 for a small mid-infrared sample, an RPD of 1.74 based on Sentinel-2, and an R^2 of 0.72 for combined GEDI and Sentinel data. Because the metrics and samples differ, the numbers alone cannot determine which performs better. Color and reflectance narrow down where laboratory analysis is needed, while field samples anchor estimates made from the sky to the ground.

Research has emerged that seeks to predict the composition of soil microbial communities through machine learning. In a 2026 Scientific Reports paper, the highest R^2, approximately 0.57, came from predictions at the phylum level, not the species level. Accuracy declined as the classification became more granular, such as at the family and genus levels, and no species-level model was presented. The limitation remains that microbial relationships learned in one field do not transfer unchanged to different soils and seasons. Performance must not be inflated by switching classification levels.

Soil carbon is traded as credits in the market, but its measured value varies by method. Remote sensing can be used to track changes across broad areas and determine sampling locations. Verra's international certification standard VM0042 v2.1 does not permit changes in soil organic carbon to be quantified through remote sensing alone and requires field measurements. Rather than portraying algorithms and laboratory analysis as two opposing camps that replace one another, the procedure of calibrating models with field samples and calculating uncertainty must be explained.

Underlying this dispute is the process of verification. Even if a farm says it followed midseason drainage and applied biochar, the resulting reductions will not be recognized in either the carbon-credit market or government subsidy programs unless they are measured according to international standards, recorded, and verified again. This process of monitoring, reporting, and verification is called MRV. Only when sensor measurements, satellite images, and artificial-intelligence predictions are placed within a single standardized framework do the figures acquire monetary value.

The problem is that the standardization process itself differs by country and certifying institution. Even for methane reduced in the same paddy, the amount recognized varies depending on where the paperwork is filed. Small farms in particular struggle to handle this paperwork alone and receive payment for carbon only after paying fees to companies that prepare the documents on their behalf.

Not all of this equipment remains fully functional in paddies and fields. Sensors buried in the soil for months accumulate layers of salt on their surfaces, microorganisms adhere to them in thin films, and moisture gradually corrodes their circuits. Values that were accurate at first begin to drift after six months or a year. Some farmers diligently maintain sensors for the first year or two, but then begin skipping calibration as it is pushed aside by other work that overlaps with each harvest season. Numbers still appear on the screen, but no one can say with confidence what they mean.

Sensors in the field must therefore be recalibrated periodically, and data from farms that skip this work lose credibility.

The price is substantial as well. A complete set of embedded sensors, drones, and artificial-intelligence analytics services carries an initial cost that is burdensome for small farms. Calls are therefore growing for more nearly free services that require nothing more than sending a photograph of soil color taken with a smartphone camera. The aim is to allow farmers to assess soil conditions with a phone in their hand, without costly embedded sensors. Who owns these data and under whose name they will be entered into carbon accounts also remains unresolved.

Emerging technologies seek to recreate the entirety of the underground environment inside a computer. The idea is to build a digital twin of soil that represents how microorganisms, carbon, and moisture interact and move, and then use it to run advance experiments regardless of how the climate changes 30 years from now. Some predict that if the volume of computation grows beyond the capacity of current computers, artificial intelligence drawing on the principles of quantum computing will narrow the immense range of possibilities far more quickly.

Large corporations and national research laboratories are already taking the lead in this kind of computation. Yet several more steps remain before these technologies reach the hands of farmers standing on paddy banks.

As we wash our hands at the paddy bank and stand up, we look down at the soil again. We cannot see how many tons of carbon this handful of soil holds or the price at which that figure will be entered in a company's ledger. Sensors, satellites, and models are filling the gap, but no one can yet answer with confidence how much the number displayed on the screen should be trusted.

Sources and References

1. Regenerative biochar for carbon sequestration and emerging technologies in soil organic carbon management for sustainable agriculture: A review. (2025). https://doi.org/10.54963/ia.v2i1.2335

2. International Journal of Agriculture and Environmental Sciences. (2025). Application of machine learning algorithms for predictive modeling of climate-smart agriculture, 1(1). https://doi.org/10.64137/3108-088x/ijaes-v1i1p103

3. AI for climate-smart agriculture: Precision sensing, crop intelligence, and resilient food systems. (2026). International Journal of Academic and Industrial Research Innovations Research Book Series. https://doi.org/10.62311/nesx/rb-978-81-999639-7-9

4. Adoption of deep learning driven precision agriculture for optimizing crop productivity and soil health via predictive analytics and autonomous sensing mechanisms. (2026). Preprints. https://doi.org/10.20944/preprints202512.2764.v1

5. Mapping climate smart agricultural interventions in rice cultivation: A lexicometric and systematic review of methane emissions and yield outcomes. (2026). Frontiers in Climate. https://doi.org/10.3389/fclim.2026.1779948

6. International Rice Research Institute. (2025). Southeast Asia tackles methane emissions with new low-emission rice production project. https://www.irri.org/news-and-events/news/southeast-asia-tackles-methane-emissions-new-low-emission-rice-production

7. Global Methane Hub. (2025, November 6). The Global Methane Hub announces $30 million investment to fast-track research and development on low-emission rice. https://www.globalmethanehub.org/2025/11/06/the-global-methane-hub-announces-30-million-investment-to-fast-track-research-and-development-on-low-emission-rice/

8. Artificial intelligence in soil microbiome-driven agriculture: From practical limits to a translational roadmap. (2026). Frontiers in Microbiomes. https://doi.org/10.3389/frmbi.2026.1860559

9. Soil microbiome prediction using traditional machine learning and deep learning models. (2026). Scientific Reports. https://doi.org/10.1038/s41598-026-39537-w

10. Abasiita, D. (2026, June 24). Why soil carbon markets hinge on a measurement few can agree on. Forbes. https://www.forbes.com/sites/daraabasiita/2026/06/24/why-soil-carbon-markets-hinge-on-a-measurement-few-can-agree-on/

11. Rural Development Administration. (2026). Overcoming the Climate and Energy Crises: Establishing Low-Carbon Rice Cultivation Technologies. Rural Development Administration press release.

https://www.rda.go.kr/board/board.do?boardId=farmprmninfo&prgId=day_farmprmninfoEntry&currPage=1&dataNo=100000809507

12. PatSnap Eureka. (2026). Soil carbon measurement technology 2026. PatSnap. https://www.patsnap.com/resources/blog/rd-blog/soil-carbon-measurement-technology-2026-patsnap-eureka/

6.2

Farm ESG Management and Community Connections

Standing on a paddy bank before sunrise, it is still dark. This paddy has been farmed for more than 30 years, and checking the water level first thing every morning has long been a habit. Today, however, the phone in the farmer's pocket sounds first. The rice processor that purchases the crop has sent a text message asking the farmer, beginning this year, to report the variety cultivated, the dates when the paddy was drained and reflooded, and the timing and amount of fertilizer applied.

This text message originated in a meeting room at corporate headquarters. A company that sells packaged meals is not responsible only for the water discharged when it washes the meal containers. It must also record in its own ledger the carbon emitted by the paddy where the rice inside those containers was grown. Shareholders and investors are pressing the company to honor its carbon-neutrality pledge. The industry calls these Scope 3 emissions. Between 70 and 90 percent of a food company's carbon footprint falls under this category, and the majority of it comes from the farms that grow its ingredients.

Even after solar panels are installed on factory roofs and delivery trucks are replaced with electric vehicles, the largest figures remain in fields and paddies. This is why a company must obtain farmers' records before it can complete its own.

Nor is the situation easy for procurement officers. To meet the target set by headquarters, they must collect the same form from hundreds of contracted production sites, yet in some years nearly half do not respond at all.

The required data are specific. According to guidance compiled in May 2026 by the U.S. agricultural information company FarmRaise, buyers ask for photographs stamped with both location and time, whether cover crops were planted, how the paddy was tilled, and what type of fertilizer was applied, when, and in what quantity. They require an unbroken record beginning with a single plot of paddy and continuing through purchase and processing to the company's final report. The European Union's Corporate Sustainability Reporting Directive, which established such standards, was originally designed for large companies.

Yet the person actually filling in the blanks is a farmer over 60 who tends ten majigi of paddy alone. With hands that can barely exchange text messages on a smartphone, the farmer must transfer satellite coordinates and fertilization dates into a table. In the past, delivering rice to an agricultural cooperative and receiving a stamp completed the year's paperwork. Now that stamp is only the beginning.

For small farms, supply-chain disclosure forms are a heavy burden. Following the Omnibus I proposal in 2025, the European Union finalized a framework in 2026 that reduces sustainability reporting and due-diligence obligations. It included a direction that small companies in the value chain should not be asked for information beyond the statutory voluntary standard. Between the year the proposal was introduced and the year the framework was finalized, farmers continued recording irrigation dates in the same notebooks.

Korea's Financial Services Commission announced its final sustainability disclosure roadmap in July 2026. Mandatory disclosure will be phased in for companies with assets of at least 10 trillion won in 2028 and those with assets of at least 5 trillion won in 2029. Scope 3 will apply beginning in 2031 to companies with assets of at least 10 trillion won. The point at which farms must submit supply-chain data will vary according to the size of the companies they do business with and the scope of disclosure.

The agricultural AI agent unveiled by the Rural Development Administration and NAVER Cloud in 2025 assisted with farming consultations and planning for people returning to farming, and recommended approximately 1,700 educational videos tailored to users. It was announced that a pest and disease image feature would be added in the first half of 2026. It cannot be written in the present tense as already available without confirming whether it was actually launched. Functions that automatically prepare farming logs and disclosure forms were likewise not part of the service unveiled at the time, but a possible direction for future expansion.

To link sensor readings to corporate sustainability disclosures, the data must be organized according to international standards. GRI 305 and GRI 101 are topical standards and disclosure items covering emissions and biodiversity, respectively. They cannot be described as two forms, Form 305 and Form 101, that are filled in automatically. The ESG Store of Taiwan's Ministry of Agriculture is a marketplace that connects corporate investment demand with emissions-reduction and adaptation projects in agriculture. Measurement standards, verification, and contracts are still required before sensor readings can be converted directly into investment.

Taiwan's New Golden Corridor project seeks to reduce groundwater use and land subsidence along the high-speed rail corridor in Changhua and Yunlin. It has pursued the planting of miscellaneous grains and dry-field crops in place of water-intensive crops and the expansion of water-saving irrigation. Its location and program are distinct from Kaohsiung City's agricultural information system. Pumping less groundwater from paddies along the rail corridor also slows the rate at which the ground beneath the tracks subsides.

Water management is not the only practice that can be quantified. Sod culture, which deliberately leaves some ground vegetation on the paddy floor; no-till farming, which uses the plow only shallowly; and returning harvested rice straw to the paddy instead of burning it all sequester carbon in the soil. This soil carbon is called Yellow Carbon, while carbon held in trees and crops is called Green Carbon, and carbon held in tidal flats and the ocean is called Blue Carbon. Though the names distinguish them by color, the work done in paddies consists of turning rice straw into the soil instead of burning it.

Demonstration projects have shown that paddies properly following intermittent irrigation can reduce methane emissions by more than 30 percent. In some rural communities in Taiwan, corporate funding has supported public meal programs for older residents.

This trend has barriers as well. Farms able to afford the initial cost of installing sensors and servers are generally those that have already reached a certain scale. The small farms that need preferential interest rates and premiums most urgently are more likely to be filtered out at the threshold.

Boomitra, which operates in India, uses satellites and artificial intelligence to estimate changes in soil carbon over broad areas of farmland. Verra VM0042 does not permit changes in soil organic carbon to be quantified through remote sensing alone. Field samples, depth criteria, model calibration, and independent verification are all required. Satellites do not replace field sampling; they first identify which points in which fields should be sampled.

This does not mean that artificial intelligence solves everything. The TRACE-RICE project, which demonstrated rice traceability in Portugal, developed an application that connects records from field to table through blockchain. Yet the researchers identified people, not technology, as the obstacle. Older farmers had more difficulty using smartphone applications, and consumers scanned QR codes only 174 times in the first demonstration in 2023. Without sustained training and assistance, the record breaks down from the moment a QR code must be attached.

Even so, the project added three more monitoring areas in 2024 and included Caravela, an indigenous Portuguese variety. A record said to extend from field to table becomes empty from the beginning the moment one person in the field stops participating.

The school meal platform supported by the TaiwanICDF manages menus, inventory, and procurement information. It is a system designed to display on one screen what and how much is delivered to each school. A blockchain-based recall network that can trace ingredients back to fields within seconds is not a confirmed function of this project. Tracing school-meal ingredients to farms would require producer codes, distribution histories, and recall procedures extending beyond the platform.

When technology enters a village, new disputes enter with it. Consider a village where 20 majigi of paddy share a single reservoir. Multiple demonstrations have already confirmed that coordinating intermittent irrigation, draining and filling paddies for set numbers of days at the reservoir level, can substantially reduce methane emissions. The problem comes next. A company's subsidy is calculated according to the amount of water saved by the reservoir as a whole and is paid to the village in a lump sum.

The paddies upstream required actual work to drain and refill them, while the paddies downstream merely received water later, yet their owners seek an equal share of the money.

The same applies to shared dryers and combines. The machine passes through paddies belonging to several households in one day, but the fuel savings are recorded under the name of the one household that operated the screen that day. If only that household receives a stamp on the preferential-rate paperwork at the next settlement, the other households remain farms with no documented performance despite having shared the machine.

The scene at the village hall on settlement day therefore differs from the past. The village head opens a laptop and projects the screen onto the wall. A graph of the reservoir water level appears beside the paddy area recorded for each household. Some nod as they look at the graph, while others point to their paddy identification number and ask for the calculation to be done again. Money once divided according to visual estimates and long-standing personal bonds is now allocated by numbers on a screen. Some people feel those numbers are wrong, but few in the village know how to identify where and how they are wrong.

Drones and satellites cannot capture only one plot of paddy in isolation in the first place. A neighboring paddy that did not enroll appears in the same image and is combined with the data for the entire village. Biodiversity records raise a similar issue. One automaker works with farms to mark locations frequently visited by the endangered leopard cat as hazard zones in its navigation system. Yet a farmer beside a paddy where leopard cats are photographed unusually often may not welcome the news. The concern is that monitoring and regulation may extend to that farmer's land.

This is the background to the concept known as farm data sovereignty: the numbers generated by my paddy belong to me. A study published in the international journal Intelligent Agriculture in 2025 observed that introducing artificial intelligence into agriculture can be sustained only when the process itself incorporates ethical responsibility and user-centered design. No rules yet exist anywhere for deciding how and to whom this right should be allocated at the village level.

At least the blueprint for reducing this burden already exists. The structure would have an app record the farmer's action at the very moment the farmer presses the irrigation button or opens a fertilizer bag, then automatically convert it into 305 and 101 forms without human intervention. The text message arriving at dawn may someday be answered by a sensor instead of a person. Until that day, however, the hand that opens the attachment and fills in the blanks still belongs to the paddy owner.

A farm cannot do this alone. The sensor manufacturer, the financing bank, the government that sets standards, the food company receiving the data, and the farmer who goes to the paddy and checks the sluice must all work together for the system to function at all. If any one party pursues only its own interests, no amount of effort by the others will make it possible to handle even a single sheet of paperwork that has reached the paddy bank.

The question that remains to the end is ultimately money. A preferential interest rate is a good program. But it only reduces interest on borrowed money; it does not raise the price paid for rice. Intermittent irrigation requires checking the paddy-water level several times a day, and spreading biochar, a charcoal-like soil amendment, adds a task that did not exist before. According to a study by a team of Indian researchers published in an international journal in 2025, biochar is an inexpensive and easy-to-handle means of sequestering carbon in soil.

Yet inexpensive refers only to the cost of the material, not to the wages of the person who spreads it evenly across the paddy.

When the settlement statement arrives in the fall, however, the price of rice produced with all this labor often differs little from that of rice from the neighboring paddy where nothing was done. In the U.S. soybean market, some regions reportedly pay a premium of $0.75 to $1.25 per bushel, approximately 27 kilograms, for certain popular varieties with favorable oil content. Farmers need such a premium to plant the same soybeans the next year.

If the price remains unchanged, they may stamp the loan documents and pose for a photograph with the official in the first year, but from the second year, when oversight becomes less frequent, they quietly begin abandoning the most labor-intensive work.

The hands that tend the sluice, check the reservoir water level, and attach QR codes all belong to the same person. In the fall, the farmer stands again on the same paddy bank and compares two lines on the settlement statement: the price per kilogram from the paddy that was drained and refilled according to a careful count of days, and the price per kilogram from the neighboring paddy where nothing was done. If the two figures are the same, can this farmer find a reason to repeat the work next year?

Sources and References

1. Sangwa, S., & Mutabazi, P. (2025). Artificial intelligence in agriculture: Ethical stewardship, responsible innovation, and governance for sustainable food systems. Intelligent Agriculture, 1(2), 25-45.
https://doi.org/10.54963/ia.v1i2.1684

2. Khatri, A., Kumar, K., & Thakur, I. S. (2025). Regenerative biochar for carbon sequestration and emerging technologies in soil organic carbon management for sustainable agriculture: A review. Intelligent Agriculture, 2(1).
https://doi.org/10.54963/ia.v2i1.2335

3. Gonçalves, C., Fernandes, J., & Brites, C. (2025). Blockchain-enabled traceability in the rice supply chain: Insights from the TRACE-RICE project. Foods, 14(21), 3711.
https://doi.org/10.3390/foods14213711

4. Rural Development Administration. (2025). Agricultural Artificial Intelligence (AI) Agent Unveiled in Partnership with NAVER Cloud. Rural Development Administration press release.
https://www.rda.go.kr/board/board.do?mode=view&prgId=day_farmprmninfoEntry&dataNo=100000805404

5. Financial Services Commission. (2026). Final Roadmap for Sustainability Disclosures.
https://www.fsc.go.kr/eng/pr010101/87276

6. Talkington, I. (2026). Supply chain disclosure: What Scope 3 buyers are asking ag producers to prove. FarmRaise.
https://www.farmraise.com/blog/2026-supply-chain-disclosure-scope-3-buyers-ag-producers

7. Key ESG. (2026). CSRD food & agriculture sector report. Key ESG.
https://resources.keyesg.com/hubfs/Whitepapers/CSRD%20Sector%20reports/KEY%20ESG%20CSRD%20Food%20&%20Agriculture%20Sector%20Report.pdf

8. Verra. VM0042 v2.1 Frequently Asked Questions: Soil Carbon Measurement Requirements.
https://verra.org/methodologies-main/frequently-asked-questions-vm0042/

9. Committee on World Food Security. (2026). Harnessing AI, digitalization and data governance for food security and nutrition: Morning session [Video]. YouTube.
https://www.youtube.com/watch?v=uQ-6ZL9UtvU

6.3

Global Food Security and the Future of Agriculture's Digital Transformation

When a bulk carrier docks at the grain terminal in Incheon's South Port, wheat, corn, and feed ingredients are unloaded by crane. Korea's grain self-sufficiency rate averaged 19.2 percent from 2023 through 2025, and one projection indicated that it could decline to 16.8 percent by 2035. The rates of 1.5 percent for wheat and 4.3 percent for corn are based on 2024. An average and figures from a specific year must not be presented as though they refer to the same point in time. The structure through which prices and transportation disruptions at import ports spread to domestic food and feed costs begins with this low self-sufficiency rate.

Even a one-day delay in a ship's arrival at this port shakes domestic prices. Yet when grain prices truly surge, Korea is not the first country where people go hungry. Countries where conflict continues, and countries unable to produce their own grain and therefore directly exposed to purchase prices, are hit first. In a 2026 report, the World Bank projected that the international food price index would rise 2.5 percent that year. It warned that this forecast could easily be exceeded if instability in the Middle East, extreme weather, and export restrictions imposed by multiple countries holding grain at their ports all coincide.

According to the 2026 global food security report issued by the Food and Agriculture Organization of the United Nations, 645 million people experienced hunger in 2025. This represented 7.8 percent of the world's population and was 43 million fewer than at the peak in 2022. The number of people unable to obtain adequate meals even if they did not experience hunger was 2.1 billion, down 86 million from one year earlier. Even so, 2.7 billion people could not afford a healthy diet. The cost of a healthy diet per person per day rose to $4.28 in 2025.

This figure is based on purchasing power parity (PPP), which adjusts different national price levels to a common standard, and is markedly higher than $2.94 in 2017. The Food and Agriculture Organization titled this report “Hunger Is Not Inevitable.” The title reflects the fact that the number of people experiencing hunger declined for three consecutive years. Yet the pace remains far too slow to meet the 2030 target.

Behind each of these statistics are people. When grain prices rise in markets across the Sahel, mothers reduce their own portions first. They leave their children's bowls unchanged and retain only broth in their own. In some refugee camps across borders, rations have been cut to one meal a day. School meals have been the final barrier preventing hunger, and when budgets shrink, that barrier is the first to be lowered. These scenes do not appear in statistical tables. What appears is only the total of 645 million.

Behind these figures are two kinds of farmers. On one side are farm owners whose fields are frequently scanned by satellites. Artificial intelligence reads multispectral images taken by drones together with values from sensors embedded in the soil, forecasting next month's yields and the risk of pests and diseases. Artificial intelligence divides the land into small sections and tells them when and how much to irrigate and which plots need more fertilizer, a practice known as variable-rate application.

When grain prices begin to fluctuate, these farm owners have the flexibility to lock in prices on the futures market or move planting dates forward.

On the other side are smallholders in Africa and South Asia. Smallholders account for 80 percent of all farmers in developing countries, yet often the only device in their hands is a feature phone with a cracked screen. In some places, the cost of a few days of mobile data is equivalent to the cost of using a diagnostic app once. A study published in 2026 identified unstable internet access, unaffordable device prices, low digital literacy, and frequent power outages as factors creating this gap.

Artificial intelligence has the power to change agriculture, but critics note that this power flows first toward farms that already possess resources.

Even if a smallholder at the edge of a field in Kenya or Rwanda photographs a diseased leaf and tries to upload it to a diagnostic app, the image will not be transmitted when the connection drops, and only a loading icon will spin on the screen. An Iowa combine, by contrast, uploads yield data continuously to the cloud even as it moves. Satellite images taken on heavily overcast days may fail to read the ground properly, causing forecasts to miss the mark, but it is still well-resourced farms that have the personnel and equipment to correct those errors.

Artificial-intelligence models are generally trained on images of large fields in the United States or Europe. As a result, they often produce incorrect answers when shown images of African fields where several crops are mixed together on small plots.

Language is another obstacle. Most large language models are trained first in languages with many users, such as English, Chinese, and Spanish. Languages such as Swahili and Hausa, which are spoken by tens of millions of people but leave relatively little written material on the internet, are pushed behind. When a Kenyan smallholder asks about a pest or disease in Swahili, the model may provide an irrelevant answer or fail to understand the question at all. The more data a language has, the more accurate the model becomes, and the more accurate it becomes, the more speakers of that language gather around it.

Languages with little data remain outside this cycle.

Around the time of the foreign-exchange crisis, several Korean seed companies passed into foreign ownership. Seoul Seed, Hungnong Seed, and Jungang Seed were sold one after another, and Cheongwon Seed, then known as the industry's seventh-largest company, was also acquired by a foreign company. The movement did not proceed in only one direction. In 2012, FarmHannong repurchased some of Monsanto Korea's seed-business rights. To understand the risk of seed and agricultural-machinery data being concentrated on the servers of a handful of companies, both directions, sale and recovery, must be considered together.

The Indian government launched Phase 1 of Bharat Vistaar on February 17, 2026. This system connects agricultural knowledge and government services through AI, while AgriStack's UFSI exchanges data subject to authorization and farmers' consent. Being connected does not mean that anyone can retrieve all the data. CIMMYT provides germplasm free of charge or at minimal cost for research, breeding, and educational purposes. Nonexclusive or quasi-exclusive licenses are used to commercialize finished seed varieties.

Reporting on investment in Korean agri-food technology startups in 2026 described an increase of 171 percent based on figures before rounding. Calculating only from the $97 million and $253 million reported in the same article gives approximately 161 percent, so the unrounded amounts must be examined. According to AgFunder, global agri-food technology investment totaled $16.2 billion in 2025. The agricultural AI market and total investment in agri-food technology are two markets with different boundaries. Before comparing the figures, it is necessary to determine whether they measure the same scope.

Over the next five years, satellite imagery, on-device models, local-language voice services, and agricultural-machinery rental are highly likely to expand. Yet this is the author's projection, not an established current fact. The point at which farms in middle-income countries receive more accurate warnings and resistant seed varieties reach more smallholders depends on communications networks, variety registration, and seed supply chains. The expectation that public infrastructure resembling Bharat Vistaar will emerge in multiple countries is likewise presented as a conditional scenario.

Power grids, communications towers, and seed distribution networks change more slowly than artificial-intelligence models. The concentration of data and patents among a handful of companies and the disruption of models by unprecedented extreme weather are likewise unlikely to disappear in a short period. No date has yet been set for farm data sovereignty to be enshrined in an international treaty or for quantum computers to calculate water and carbon across the entire planet. While inexpensive chips are installed one by one along paddy banks, the servers collecting their data may grow even faster.

The cranes at Incheon port continue to lift wheat today. By the time that wheat becomes noodles and bread and reaches the table, the smallholder holding a feature phone with a cracked screen has not been included anywhere in this calculation. When artificial intelligence uses cheaper satellite images and faster chips to read every corner of cornfields and wheat fields within five years, who will be the first to receive the results?

Sources and References

1. Sangwa, S., & Mutabazi, P. (2025). Artificial Intelligence in Agriculture: Ethical Stewardship, Responsible Innovation, and Governance for Sustainable Food Systems. Intelligent Agriculture, 1(2). https://doi.org/10.54963/ia.v1i2.1684

2. Press Information Bureau, Government of India. (2026). Rollout of BHARAT-VISTAAR Platform. https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2239788

3. SHARE GENES. (2026). Predictive Breeding: Bridging Genetics, Artificial Intelligence and Agriculture | Dr. B. M. Prasanna [Video]. YouTube. https://www.youtube.com/watch?v=FWcxeEC90JA

4. Committee on World Food Security. (2026). Morning Session: Harnessing AI, Digitalization and Data Governance for Food Security and Nutrition [Video]. YouTube. https://www.youtube.com/watch?v=uQ-6ZL9UtvU

5. Food and Agriculture Organization of the United Nations. (2026). SOFI 2026: Hunger not inevitable, FAO says as global hunger report shows progress. FAO Newsroom. https://www.fao.org/newsroom/detail/sofi-2026--hunger-not-inevitable--fao-says-as-global-hunger-report-shows-progress/en

6. World Bank. (2026). When risks stack up: Threats to global food markets in 2026. World Bank Blogs. https://blogs.worldbank.org/en/developmenttalk/when-risks-stack-up--threats-to-global-food-markets-in-2026

7. Financial News. (2026). Korea's Grain Self-Sufficiency Rate Projected to Fall to the 16 Percent Range by 2035, Deepening Food Security Vulnerability. https://www.fnnews.com/news/202607020630507050

8. AgFunder. (2026). Global AgriFoodTech Investment Report 2026. https://agfunder.com/research/agfunder-global-agrifoodtech-investment-report-2026/

9. Phys.org. (2026). AI offers promise for agriculture, but smallholder farmers risk being left behind. https://phys.org/news/2026-06-ai-agriculture-smallholder-farmers-left.html

10. NewsPim. (2026, June 22). [Reading the Economy Through AI] When Seeds Become Hegemony: The $62 Billion Seed War Has Begun. https://www.newspim.com/news/view/20260622000833

11. NewsPim. (2026, June 23). [Reading the Economy Through AI] The Seeds Sold Off by the IMF: A 25-Year Record of Korea's Seed Sovereignty. https://www.newspim.com/news/view/20260622001057

12. NewsPim. (2026, June 27). [Reading the Economy Through AI] After K-Semiconductors Comes K-Seeds: Transform the Rural Development Administration into an Industrial Platform. https://www.newspim.com/news/view/20260626000833

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.

PayPal

https://paypal.me/kimkjkj

Scan the QR code or open the link above.