Scalable Agricultural Intelligence

AI as an Input to Agricultural Production

Johnny Kao

Independent Researcher, Tokyo, Japan

ORCID: 0009-0004-1446-6001

Working Paper

Abstract

Artificial intelligence in agriculture is often described as a substitute for labor or as a tool for improving prediction, precision, and control. Yet many agricultural systems are constrained by a different scarce resource: experienced judgment. A farm may have land, capital, equipment, workers, and demand, but still depend on a limited number of people who know what to notice, how to interpret changing conditions, and when to intervene. This paper develops the Scalable Agricultural Intelligence framework by formalizing the argument in Kao (2026a) that intelligence can be treated as an agricultural input alongside more visible inputs such as land, labor, water, energy, and capital. The term does not imply a standardized factor of production. It identifies the operational capacity to observe conditions, interpret their significance, and select useful actions.

The paper connects agricultural extension, decision-support systems, human-centered AI, agricultural expertise, and the economics of knowledge in organizations. It introduces four definitions: Agricultural Intelligence, Judgment Bottleneck, Expertise Reach, and Scalable Agricultural Intelligence. The central mechanism is that AI can reduce the amount of routine expert judgment required per unit of production, route exceptional cases toward scarce experts, lower the knowledge required at the point of entry for less experienced operators, and partially separate expertise from physical location. A lightweight formulation expresses expertise reach as a capacity relationship rather than a calibrated production function.

Evidence from digital agricultural extension, autonomous greenhouse experiments, human-centered agricultural AI, and recent workplace AI research supports components of this mechanism while also showing its limits. Digital advice can expand knowledge reach without automatically raising output; agricultural skill contains embodied and contextual elements that are not readily codified; and AI benefits may be largest where routine decisions can be supported while humans retain responsibility for unusual conditions. The paper therefore argues that AI need not replace agricultural experts to change agricultural economics. It may be sufficient for the same scarce expertise to support more production, more locations, or less experienced operators, provided that the system preserves reliable escalation to human judgment.

Keywords: agricultural AI; agricultural intelligence; expertise; decision support; agricultural extension; human-centered AI; knowledge scaling

1. Introduction

Agriculture is commonly described through visible inputs. Land, labor, water, fertilizer, machinery, energy, and capital appear in farm budgets, policy statistics, and production models. Experienced judgment is less visible, even though many agricultural operations depend on it continuously. A grower decides whether a crop is merely stressed or becoming diseased. A farm manager decides whether a weather change requires irrigation, ventilation, or no intervention. An agronomist distinguishes an unusual event from ordinary variation. These decisions are not incidental to production. In many systems they are part of the operating infrastructure that makes production possible.

This becomes economically important when judgment is scarce. Hiring an additional pair of hands is not the same as finding another person who has seen the same crop across many seasons, understands local conditions, and can distinguish a harmless anomaly from the beginning of a major problem. The distinction is especially relevant in agriculture because biological systems are dynamic, local conditions differ, and experience accumulates slowly. Legun, Burch, and Klerkx (2023), for example, show that agricultural expertise includes both technical knowledge that may be transferable and embodied expertise that is more difficult to automate or detach from practice. Digital-agriculture research likewise emphasizes that technology changes skills, work, and agricultural knowledge systems rather than simply replacing manual tasks (Klerkx, Jakku, & Labarthe, 2019).

Artificial intelligence changes this problem because some judgment can now be mediated through software. Sensors can observe conditions continuously; models can combine weather, soil, crop, and operational data; decision-support systems can suggest or select actions; and autonomous systems can execute repeated adjustments. The important economic change is therefore not limited to whether a machine performs a physical task. It also concerns where operational judgment resides, how often an expert must intervene, and how many production units a scarce expert can support.

The source concept for this paper appears in Kao (2026a), which argues that the farmer's head is part of the farm's infrastructure and that intelligence should be treated as an agricultural input. The book distinguishes a shortage of hands from a shortage of judgment and proposes that AI can matter by multiplying the reach of expertise rather than eliminating experts. This working paper formalizes that argument and positions it within adjacent literatures on agricultural extension, decision support, agricultural skill, human-centered AI, and the economics of knowledge.

The research question is: what changes economically when agricultural judgment and expertise become partially scalable through software? The contribution is integrative. Agricultural extension research already studies how information can reach more farmers; decision-support research studies how digital systems assist farm decisions; human-centered AI research emphasizes augmentation rather than full replacement; and organizational economics has long examined how scarce problem-solving knowledge is allocated across workers. The framework developed here connects these literatures at the level of agricultural production capacity. It asks how technology changes the amount of production that a unit of scarce expert judgment can support.

2. Related Literature and the Scaling Gap

2.1 Agricultural extension and the reach of knowledge

Agricultural extension has always been partly a scaling problem. One agronomist, extension worker, or experienced farmer cannot be physically present at every farm. Digital channels can relax that constraint by distributing information more cheaply and widely. Xu et al. (2023) review technology applications in agricultural extension and document the growing role of digital tools in knowledge dissemination and farmer support. Singh, Slotznick, and Stein (2023), using a cluster-randomized trial in Southern India, find that mobile extension improved farmers' recall, knowledge, and adoption of promoted practices. The intervention substantially increased informational reach, although the study did not find statistically significant effects on production or yield. That result is useful for the present framework: scaling access to knowledge is not the same as scaling effective agricultural judgment, and neither automatically translates into production outcomes.

Recent work on AI-enabled extension moves closer to the present question. High, Singh, and Nemes (2026) argue that AI in agricultural extension should be understood within broader learning systems rather than as a simple technology-transfer mechanism. Their emphasis on participation, learning, and farmer agency cautions against treating expert knowledge as a static package that can simply be copied into software. In parallel, Japan's policy environment shows why the scaling problem matters. The Ministry of Agriculture, Forestry and

Fisheries reports that Japan had about 1.036 million core agricultural workers in 2025, with an average age of 67.7; 69.6% were age 65 or older. The ministry also notes that less than 30% of farm entities had secured a successor within five years and explicitly identifies the transfer of skills and know-how as part of succession (MAFF, 2026). OECD (2024) similarly links smart agriculture in Japan with skills development and knowledge transfer.

2.2 Decision support, autonomy, and human expertise

Agricultural decision-support systems already operationalize parts of expert judgment. Rose et al. (2016) show that farmers and advisers use decision-support tools when systems are relevant, usable, credible, and compatible with practice. Ara et al. (2021) find that agricultural decision-support systems can transmit heuristics and support tactical and strategic choices, while also noting persistent problems of adoption, uncertainty, and fit with user needs. More recent work continues this trajectory. Brugler et al. (2024) identify trust, black-box behavior, workforce capability, and training as barriers to machine-learning decision support in precision agriculture. Mouratiadou et al. (2023) present the Digital Agricultural Knowledge and Information System as a knowledge-based decision-support architecture intended to integrate data, spatial context, and stakeholder needs. A recent technical review also uses the term agricultural intelligence to organize AI capabilities around perception, decision, and execution (Yu et al., 2026). That framing is complementary to, but distinct from, the production-economic question here: how software changes the amount of agricultural production that a given stock of scarce expert judgment can support.

These literatures show that software can participate in agricultural decision-making, but they do not generally treat expert judgment itself as a scarce production input whose effective supply can be expanded. Human-centered AI work provides another part of the bridge. Holzinger et al. (2024) argue for Agriculture 5.0 systems that augment human capabilities and retain human oversight, reflecting the practical reality that not every agricultural decision is suitable for full automation. Singh and Singh (2026) similarly develop a human-in-the-loop expert system for data-scarce agriculture that combines explicit agronomic rules with constrained machine learning. The implication is not that humans disappear, but that the division of judgment between software and people can be redesigned.

2.3 Expertise as an organizational bottleneck

The organizational-economics literature makes the scaling mechanism clearer. Garicano (2000) models production as a hierarchy of problem solving: common problems are handled by workers with routine knowledge, while rarer or harder problems are escalated to specialized experts. Information technology changes organizational design when it changes the cost of acquiring, communicating, or applying knowledge. The analogy to agriculture is direct but not complete. A farm generates a stream of recurring operational decisions, many of which are ordinary and some of which require scarce experience. If software can reliably absorb more of the ordinary stream, a limited pool of expert judgment can be concentrated on exceptions.

Recent evidence from outside agriculture supports the plausibility of such knowledge scaling. Brynjolfsson, Li, and Raymond (2025) study a generative-AI assistant used by 5,172 customer-support agents and find a 15% average productivity increase, with the largest gains among less experienced and lower-skilled workers. They provide evidence that AI can disseminate behaviors associated with more productive agents. Agriculture differs sharply from customer support, but the mechanism is relevant: software can sometimes make accumulated expert practices available to people who have not yet accumulated the same experience themselves.

Table 1. Adjacent literatures and the remaining scaling question

Literature	Primary question	What it establishes	Remaining question for this paper
Agricultural extension	How can knowledge reach more farmers?	Digital delivery can increase reach, recall, and knowledge.	Can operational judgment itself become less locally scarce?
Decision support	How can software improve farm decisions?	Models and tools can support tactical and strategic choices.	How does support change the amount of production one expert can supervise?
Human-centered AI	How should humans and AI share work?	Augmentation and escalation remain important where full automation is unreliable.	What is the production-economic effect of reallocating routine judgment?
Agricultural expertise	What parts of agricultural skill are technical, embodied, or contextual?	Not all expertise is codifiable or automatable.	How much expertise must remain scarce for the economics to change?
Knowledge hierarchies	How is scarce problem-solving knowledge allocated in organizations?	Routine problems can be separated from exceptional problems.	How does this logic apply to biological production across farms and locations?

The gap is therefore narrower than a claim that AI can assist farmers or distribute knowledge. The missing analytical object is the scalability of agricultural intelligence as an input to production: how much effective agricultural operation can be supported by a given stock of scarce expert judgment when some observation, interpretation, and routine decision-making moves into software.

3. The Scalable Agricultural Intelligence Framework

Definition 1 (Agricultural Intelligence). the operational capacity to observe agricultural conditions, interpret what those conditions are likely to mean, and select a useful action for a specified production context. It may be supplied by experienced people, formal rules, models, software, or combinations of these sources.

This definition is intentionally narrower than general intelligence and broader than information. A weather forecast is information; deciding whether that forecast requires changing irrigation, ventilation, harvest timing, or nothing at all is an exercise of agricultural intelligence. The definition also avoids treating AI systems as autonomous experts by default. A system may contribute to one stage of observation, interpretation, or action selection without reliably performing all three.

Definition 2 (Judgment Bottleneck). a production constraint that arises when the supply of sufficiently reliable agricultural judgment is too limited to support the desired scale, location, continuity, or complexity of an operation, even when other necessary inputs are available.

A judgment bottleneck differs from a shortage of hands. A farm may be able to hire workers to harvest, move materials, or perform routine tasks while still lacking enough people who know what should happen next. The distinction also differs from an information shortage. More sensor data does not remove a judgment bottleneck if nobody or no system can reliably interpret the data and translate it into action.

Definition 3 (Expertise Reach). the amount of agricultural production that a unit of scarce expert judgment can effectively supervise within a specified period while maintaining an acceptable level of performance and risk.

Expertise Reach is context dependent. It may be expressed as hectares, greenhouses, herds, production sites, crop cycles, or other operating units. It is not a universal productivity ratio. The relevant measure depends on the production system and on what decisions still require expert intervention.

Definition 4 (Scalable Agricultural Intelligence). a condition in which technology increases Expertise Reach or lowers the expert knowledge required at the point of routine operation without requiring a proportionate increase in scarce human expertise.

The definition does not require expertise to be fully automated. It requires only that effective agricultural judgment become less tightly coupled to one person's continuous presence. This can occur through decision support, remote monitoring, automated control, exception routing, codified operating rules, or combinations of these mechanisms.

3.1 A lightweight capacity relation

Let B_J denote the amount of scarce expert-judgment capacity available during a period, and let j_i(T) denote the amount of that capacity required by production unit i under technology state T. A necessary capacity condition for supervising n production units is:

$$\Sigma_{i=1}^{n} j_i(T) \le B_J$$

If a technology reduces the routine expert-judgment requirement for each unit while keeping exception handling within acceptable bounds, the same B_J can support a larger n. For a set of broadly comparable units with average expert-judgment demand $\bar{j}(T)$, Expertise Reach can be represented conceptually as:

$$ER(T) = B_J / \bar{j}(T)$$

These expressions are accounting relationships, not empirically calibrated production functions. They do not imply that expert judgment is perfectly divisible, that production units are homogeneous, or that lower human involvement is always desirable. Their purpose is to make the scaling mechanism explicit: AI changes production capacity when it changes how much scarce expert judgment each operating unit consumes.

An illustrative way to read the relation is to hold expert capacity fixed and change only the amount of routine judgment each production unit consumes. If software absorbs routine climate, irrigation, and monitoring decisions across several comparable greenhouses, a senior grower may be able to supervise more sites without any increase in personal working time. This example is deliberately qualitative: the framework does not assume that expert minutes are perfectly substitutable across sites, and empirical applications must account for exception frequency, task switching, and site heterogeneity.

3.2 The entry knowledge threshold

A second mechanism concerns new or less experienced operators. Let K_min(T) denote the minimum local operating knowledge required for a person to manage routine conditions safely under technology state T. If decision support, monitoring, and automated control reliably handle part of the ordinary decision load, the required starting threshold may fall even though the operator still needs training and access to escalation. Conceptually:

$$K_{min}(T_1) < K_{min}(T_0)\ when\ supported\ routine\ judgment\ is\ reliable$$

This is not equivalent to making a novice an expert. It means that a novice may be able to begin productive operation before personally accumulating every routine heuristic that an experienced operator previously had to carry in memory. The distinction is central to agricultural succession, where years of seasonal experience cannot be reproduced quickly.

4. How Agricultural Intelligence Becomes Scalable

4.1 Routine judgment moves before expertise disappears

The first mechanism is routine-decision absorption. Agricultural operations generate many repeated decisions that are individually small but collectively consume expert time: climate set-point adjustments, irrigation timing, anomaly checks, feeding decisions, equipment monitoring, or routine diagnosis. Decision-support and autonomous-control systems can reduce the frequency with which a person must make each ordinary decision. The economic effect is not simply labor saving. It is a reduction in scarce expert-judgment demand per operating unit.

The autonomous-greenhouse literature provides a concrete example. Hemming et al. (2020) report the first international Autonomous Greenhouse Challenge, in which AI teams remotely managed cucumber crops using greenhouse sensor data and control systems; one AI-controlled compartment outperformed the manually operated expert reference on the competition's combined criteria. In the fourth challenge, Wageningen University & Research reported that algorithms controlled lighting, heating, CO_2, water, plant density, and harvest timing during the cultivation period, with the winning team achieving the highest profit (Wageningen University & Research, 2025). These competitions do not prove that commercial greenhouse expertise can be fully automated, but they demonstrate that a substantial stream of crop-management decisions can be mediated remotely through software.

4.2 Exception routing increases Expertise Reach

The second mechanism is exception routing. If software handles normal conditions but reliably identifies situations that exceed its competence, experts can spend more time on genuinely unusual cases. This resembles Garicano's knowledge hierarchy: common problems are resolved at a lower level, while difficult problems are escalated. In agriculture the lower level may include sensors, control rules, predictive models, or local operators assisted by software; the higher level may be a senior grower, veterinarian, agronomist, or specialist who can supervise several sites rather than remaining continuously attached to one.

This mechanism should not be confused with governance-oriented human-in-the-loop design. The present question is production capacity: how much agricultural activity can scarce expertise support? Reliable escalation is nevertheless a necessary operating condition. A system that reduces routine workload but fails to recognize abnormal conditions can increase apparent Expertise Reach while increasing hidden operational risk.

4.3 Software can lower the starting knowledge requirement

The third mechanism is assistance to less experienced operators. Brynjolfsson et al. (2025) show in a non-agricultural setting that AI support disproportionately improved the performance of less experienced workers and appeared to disseminate practices associated with stronger performers. Agricultural extension research shows a related but more modest effect: digital channels can increase knowledge and adoption, but the effect on output is not automatic (Singh et al., 2023). Together these findings support a restrained proposition. AI can sometimes reduce how much routine knowledge must be personally accumulated before an operator becomes useful, but production gains depend on the quality of advice, the surrounding system, and the operator's ability to act on it.

In aging agricultural systems, this mechanism also has a continuity dimension. When operating knowledge is concentrated in an older grower, succession requires more than replacing labor; it requires transferring enough judgment for the farm to keep functioning. AI-mediated records, decision support, and automated control can lower the amount of expertise that must be embodied in the incoming operator at the moment of succession, while experienced human judgment remains available for exceptions and learning. In this sense, scalable agricultural intelligence can act as succession infrastructure rather than as a substitute for the successor (Kao, 2026a).

4.4 Expertise can become partially portable

The fourth mechanism is geographic portability. Agricultural expertise has traditionally been expensive to separate from place because observation and intervention occurred on site. Sensors, cameras, connectivity, remote-control systems, and digital records can make part of the farm observable from elsewhere. This does not remove geography. Repairs, physical inspection, local knowledge, and many embodied tasks remain site-specific. It can, however, reduce how often scarce judgment must travel physically to the production site.

Table 2. Mechanisms of scalable agricultural intelligence

Mechanism	What technology changes	Production effect	Primary failure mode
Routine-decision absorption	Software performs or recommends recurring decisions.	Less expert time required per production unit.	Automation of conditions outside the model's reliable domain.
Exception routing	Normal cases are filtered; unusual cases are escalated.	Scarce experts can supervise more units.	Missed anomalies or poor escalation thresholds.
Lower entry threshold	Routine heuristics become available through software support.	Less experienced operators can begin useful work sooner.	Over-reliance before sufficient local understanding develops.
Geographic portability	Observation and some control become remotely accessible.	Expertise can support additional locations.	Loss of embodied, sensory, or place-specific information.

5. Evidence and Operational Interpretation

No single existing study validates the full Scalable Agricultural Intelligence framework. The relevant evidence instead supports separate links in the mechanism. The purpose of this section is therefore illustrative rather than confirmatory.

5.1 Extension: reach can scale before outcomes do

The Southern India extension experiment is useful because it separates reach and knowledge from production. Singh et al. (2023) find that mobile delivery improved recall, knowledge, and adoption relative to video-only extension, while production and yield effects were not statistically significant. From the present perspective, this demonstrates that one dimension of agricultural intelligence - access to applicable knowledge - can become more scalable without automatically removing every other constraint on output. Scalable intelligence is therefore not a claim that information alone produces yield.

5.2 Autonomous greenhouse control: judgment can move into software

The Wageningen greenhouse experiments demonstrate a stronger form of scaling because the systems participate directly in operational decisions. Hemming et al. (2020) report remote algorithmic control of greenhouse climate and irrigation against an expert-grower reference, while subsequent competitions expanded the scope of autonomous cultivation. These studies show that agricultural judgment can be represented and exercised through a combination of sensing, models, optimization, and remote interfaces. They do not establish that the full knowledge of expert growers has been copied. Instead, they show that enough of the routine decision stream can move into software for the organization of expertise to change.

5.3 Japan: succession makes expertise scarcity visible

Japan illustrates why the distinction between labor and judgment matters. MAFF (2026) reports a declining and aging core agricultural workforce and identifies the transfer of techniques and know-how as part of farm succession. An aging farmer leaving production can therefore remove more than one unit of labor: the exit can also remove accumulated knowledge about fields, crops, weather responses, customers, and operating routines. This makes agricultural succession a natural setting for testing whether AI functions as succession infrastructure. The empirical question is not whether a novice receives more information, but whether a farm can continue or transition with less dependence on one individual's accumulated memory.

5.4 Expertise is only partially codifiable

A credible scaling framework must also explain what does not scale. Legun et al. (2023) distinguish dimensions of agricultural expertise and show why embodied, contextual, and socially recognized skill complicates simple automation narratives. Holzinger et al. (2024) similarly argue for human-centered agricultural AI that augments rather than removes human capability. Singh and Singh (2026) use a human-in-the-loop design precisely because data-scarce agricultural settings require transparent expert rules, user control, and conservative system boundaries. These studies support a limited form of scalability: technology may expand the reach of some judgment while leaving tacit, sensory, unusual, or responsibility-laden decisions with people.

Table 3. Evidence mapped to the scalable-intelligence mechanism

Evidence	Observed result	Framework interpretation	What it does not establish
Digital extension RCT (Singh et al., 2023)	Higher knowledge and practice adoption; no clear yield effect.	Knowledge reach can scale independently of production outcomes.	That digital advice removes all production constraints.
Autonomous Greenhouse Challenge (Hemming et al., 2020; WUR, 2025)	Remote algorithms controlled repeated crop-management decisions and achieved competitive economic results.	A significant share of routine judgment can move into software.	That commercial grower expertise is fully replaceable.
Generative AI at Work (Brynjolfsson et al., 2025)	Largest productivity gains accrued to less experienced workers.	AI can transmit high-performing practices and lower effective experience gaps.	That the same effect magnitude applies to agriculture.
Agricultural expertise studies (Legun et al., 2023)	Expertise includes embodied and contextual dimensions.	Scalability is partial; exception handling and local knowledge remain important.	That tacit knowledge can be fully codified.
Japan workforce and succession data (MAFF, 2026)	Aging workforce, weak successor coverage, explicit need to transfer know-how.	Expert judgment can be a continuity constraint, not merely a labor cost.	That AI by itself solves succession.

6. Implications for Agricultural Organization

6.1 The relevant unit is not jobs replaced but production supported

Many discussions of agricultural automation ask how many workers a technology removes. Scalable Agricultural Intelligence suggests a different metric when judgment is the scarce resource: how much additional production can the same expert population support? A technology can have little effect on headcount yet materially change capacity if one greenhouse manager can supervise several facilities, one agronomist can support more farms, or one experienced grower can mentor multiple less experienced operators through exception-based intervention.

6.2 Intelligence can be complementary to labor and capital

Treating agricultural intelligence as an input does not mean inserting a new scalar into a conventional production function. The point is operational complementarity. Land and capital can remain underused when management capacity is the binding constraint. Additional workers can remain ineffective when they lack guidance. Sensors can generate little value when observations are not converted into decisions. Intelligence combines with other inputs by determining whether and how they are used. This is why a small change in judgment capacity can have a larger effect on physical production than its apparent labor saving would suggest.

That complementarity is not free. Scaling judgment through software may require sensors, connectivity, integration with existing control systems, maintenance, software subscriptions, data management, and staff adaptation. An increase in Expertise Reach is therefore an operational gain, not automatically an economic gain. Scalable Agricultural Intelligence is commercially meaningful only when the value of additional production, continuity, or reduced expert scarcity exceeds the enabling and operating costs required to achieve it.

6.3 Scalable intelligence can relax a viability constraint without guaranteeing frontier expansion

The relationship to the Agricultural Viability Frontier is direct but conditional. If scarce expertise is binding on whether a farm can expand, continue, or operate in a location, increasing Expertise Reach can relax that constraint and potentially contribute to frontier expansion (Kao, 2026b). But Scalable Agricultural Intelligence does not imply frontier expansion by itself. A farm may still be constrained by water, energy, capital, logistics, regulation, or demand. The present framework explains one mechanism through which an expertise constraint can move; the viability framework asks whether that movement changes the economic production set.

7. Research Agenda

The framework generates empirical questions that are more specific than whether farmers adopt AI. A first research program should measure expert-judgment demand directly. In greenhouse operations, livestock systems, or large field operations, researchers could record the time senior operators spend on routine monitoring, ordinary adjustments, exception handling, and recovery from system errors before and after adoption of decision-support or autonomous-control technology. The primary outcome would be expert minutes per production unit, complemented by yield, quality, loss events, and operating cost.

Measurement should also account for cognitive switching. Expert minutes spent on intervention are not fully comparable when a senior operator alternates rapidly across sites or must reconstruct local context before acting. Field studies should therefore record direct intervention time together with monitoring time, switching frequency, reorientation time, exception burstiness, and recovery from system errors. A useful empirical measure is effective expert attention consumed per managed unit, rather than nominal minutes alone.

A second program should test Expertise Reach. Multi-site firms provide a natural setting because management span can be observed as the number of greenhouses, barns, fields, machines, or hectares supervised per senior operator. Staggered technology rollouts could support difference-in-differences or event-study designs where adoption timing is plausibly exogenous to short-term operating shocks. The key outcome is not only labor cost but whether management span expands without deterioration in performance or risk.

A third program should study the entry knowledge threshold. Researchers could compare novice operators receiving standard training with operators receiving training plus AI decision support, measuring time to independent routine operation, escalation frequency, error rates, learning retention, and performance after assistance is removed. Brynjolfsson et al. (2025) suggest that assistance may compress performance differences in other knowledge-intensive work; agriculture requires direct testing because biological variability, delayed feedback, and local context may limit transferability.

A fourth program should examine succession and continuity. In aging agricultural regions, longitudinal studies could compare farms with different levels of digital knowledge capture, remote supervision, and AI-enabled decision support. Relevant outcomes include successful management transfer, farm continuation, time required for a successor to assume responsibility, and the persistence of site-specific operating knowledge after an experienced farmer exits.

Finally, research should distinguish knowledge portability from knowledge erosion. Klerkx et al. (2019) and Legun et al. (2023) highlight changes in skills and identities associated with digital agriculture. Systems that make routine expertise portable may also reduce opportunities for new operators to develop embodied understanding. Longitudinal research should therefore test whether AI-supported operators eventually acquire deeper judgment or remain dependent on the system for decisions they would otherwise learn to make.

8. Limitations

First, agricultural intelligence is a conceptual category, not an empirically standardized input. Observation, interpretation, and action selection vary greatly across crops, production systems, and locations. The framework therefore provides a way to identify and measure judgment constraints in context; it does not assume a universal unit of intelligence.

Second, the framework does not assume that all agricultural expertise can be codified. Embodied perception, sensory cues, social knowledge, repair skills, local environmental memory, and genuinely novel situations may remain difficult to capture digitally. Scalable Agricultural Intelligence can be economically important even when only routine or moderately complex decisions become software-mediated.

Third, increasing Expertise Reach may create new dependencies and failure modes. A larger operating span can concentrate risk if many production units depend on the same model, data pipeline, connectivity layer, or remote expert. Apparent scalability can therefore reduce resilience when failures are correlated. Human-centered design, redundancy, auditability, and reliable escalation remain important even though they are not the primary subject of this paper.

Fourth, the evidence assembled here is heterogeneous and does not validate a single causal model. Extension trials, greenhouse competitions, organizational-economics evidence, and agricultural expertise studies operate at different scales and measure different outcomes. They support the plausibility of separate mechanisms rather than an empirically calibrated estimate of Scalable Agricultural Intelligence.

Fifth, more scalable judgment is not necessarily socially desirable in every application. Technologies can alter labor relations, data ownership, autonomy, skill formation, and the distribution of decision authority. The present paper focuses on production capacity and expertise constraints, not on the full ethical or political economy of agricultural digitalization.

9. Conclusion

Agriculture does not depend only on physical inputs. It also depends on people who know what to notice, what matters, and what to do next. When that judgment is scarce, it can become a production bottleneck even when land, capital, equipment, workers, and customers are available.

Scalable Agricultural Intelligence provides a way to analyze what changes when part of that judgment becomes software-mediated. The framework distinguishes agricultural intelligence from information, identifies judgment bottlenecks, defines Expertise Reach, and describes four mechanisms through which AI can expand effective judgment capacity: routine-decision absorption, exception routing, a lower entry knowledge threshold, and partial geographic portability.

The central claim is deliberately limited. AI does not need to reproduce the entire farmer, agronomist, or greenhouse manager to change agricultural economics. If it can handle enough ordinary observation and decision-making that scarce experts support more production, or that less experienced operators can begin useful work with reliable escalation, the effective supply of agricultural intelligence changes. That shift may matter as much as replacing a physical task, particularly in production systems constrained by aging expertise, management capacity, or distance.

AI Use Disclosure

OpenAI's ChatGPT was used only for literature discovery and bibliographic verification. The author independently reviewed all cited sources and remains solely responsible for the final manuscript, its arguments, source selection, and citations.

References

Ara, I., Turner, L., Harrison, M. T., Monjardino, M., deVoil, P., & Rodriguez, D. (2021). Application, adoption and opportunities for improving decision support systems in irrigated agriculture: A review. Agricultural Water Management, 257, 107161. https://doi.org/10.1016/j.agwat.2021.107161

Brugler, S., Gardezi, M., Dadkhah, A., Rizzo, D. M., Zia, A., & Clay, S. A. (2024). Improving decision support systems with machine learning: Identifying barriers to adoption. Agronomy Journal, 116(3), 1229-1236. https://doi.org/10.1002/agj2.21432

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889-942. https://doi.org/10.1093/qje/qjae044

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. Journal of Political Economy, 108(5), 874-904. https://doi.org/10.1086/317671

Hemming, S., Righini, I., de Zwart, F., & Elings, A. (2020). Remote control of greenhouse cucumber production with artificial intelligence - results from the first international autonomous greenhouse challenge. Acta Horticulturae, 1296, 987-996. https://doi.org/10.17660/ActaHortic.2020.1296.125

High, C., Singh, N., & Nemes, G. (2026). Artificial intelligence for agricultural extension: Supporting transformative learning among smallholder farmers. Journal of Development Policy and Practice, 11(1), 61-80. https://doi.org/10.1177/24551333251345224

Holzinger, A., Fister, I., Jr., Fister, I., Sr., Kaul, H.-P., & Asseng, S. (2024). Human-centered AI in smart farming: Toward Agriculture 5.0. IEEE Access, 12, 62199-62214. https://doi.org/10.1109/ACCESS.2024.3395532

Kao, J. (2026a). Farming the impossible: How AI is expanding what agriculture can do. ISBN 978-626-01-7158-2.

Kao, J. (2026b). The Agricultural Viability Frontier: When AI expands what farming can economically do. Working paper.

Klerkx, L., Jakku, E., & Labarthe, P. (2019). A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda. NJAS - Wageningen Journal of Life Sciences, 90-91, 100315. https://doi.org/10.1016/j.njas.2019.100315

Legun, K., Burch, K. A., & Klerkx, L. (2023). Can a robot be an expert? The social meaning of skill and its expression through the prospect of autonomous AgTech. Agriculture and Human Values, 40(2), 501-517. https://doi.org/10.1007/s10460-022-10388-1

Ministry of Agriculture, Forestry and Fisheries of Japan. (2026). FY2025 Annual Report on Food, Agriculture and Rural Areas. https://www.maff.go.jp/j/wpaper/w_maff/r7/index.html

Mouratiadou, I., Lemke, N., Chen, C., Wartenberg, A., Bloch, R., Donat, M., Gaiser, T., Hanike Basavegowda, D., Helming, K., Hosseini Yekani, S. A., Krull, M., Lingemann, K., Macpherson, J., Melzer, M., Nendel, C., Piorr, A., Shaaban, M., Zander, P., Weltzien, C., & Bellingrath-Kimura, S. D. (2023). The Digital Agricultural Knowledge and Information System (DAKIS): Employing digitalisation to encourage diversified and multifunctional agricultural systems. Environmental Science and Ecotechnology, 16, 100274. https://doi.org/10.1016/j.ese.2023.100274

OECD. (2024). Agricultural Policy Monitoring and Evaluation 2024: Innovation for Sustainable Productivity Growth. OECD Publishing. https://doi.org/10.1787/74da57ed-en

Rose, D. C., Sutherland, W. J., Parker, C., Lobley, M., Winter, M., Morris, C., Twining, S., Ffoulkes, C., Amano, T., & Dicks, L. V. (2016). Decision support tools for agriculture: Towards effective design and delivery. Agricultural Systems, 149, 165-174. https://doi.org/10.1016/j.agsy.2016.09.009

Singh, R., Slotznick, W., & Stein, D. (2023). Digital tools for rural agriculture extension: Impacts of mobile-based advisories on agricultural practices in Southern India. Journal of the Agricultural and Applied Economics Association, 2(1), 4-19. https://doi.org/10.1002/jaa2.42

Singh, E. B. K., & Singh, T. S. (2026). AgriField-Manipur: A human-in-the-loop knowledge-based expert system for farm-level decision support under data scarcity. Expert Systems with Applications, 324, 132522. https://doi.org/10.1016/j.eswa.2026.132522

Wageningen University & Research. (2025). Winner of the 4th Autonomous Greenhouse Challenge announced. https://www.wur.nl/en/news/winner-4th-autonomous-greenhouse-challenge-announced

Xu, Z., Adeyemi, A. E., Catalan, E., Ma, S., Kogut, A., & Guzman, C. (2023). A scoping review on technology applications in agricultural extension. PLOS ONE, 18(11), e0292877. https://doi.org/10.1371/journal.pone.0292877

Yu, S., Li, X., Zhao, S., & Liu, Q. (2026). Agricultural Intelligence: A technical review within the perception-decision-execution framework. Applied System Innovation, 9(5), 95. https://doi.org/10.3390/asi9050095