

ORIGINAL RESEARCH

Open access

Human-Centred Autonomous Farming: A Framework for Trust, Override, Explainability, and Safety in Robotic Agricultural Systems

Mina Takeda^{1*}, Haruto Suzuki¹

Abstract

Autonomous farming systems are moving rapidly from experimental prototypes into commercial and operational agricultural settings. Autonomous tractors, robotic weeders, drones, harvest-assist systems, and sensor-guided machines promise to address labour shortages, improve operational efficiency, and support more precise field management. Yet the governance arrangements needed to ensure that these systems remain safe, understandable, and accountable have not developed at the same pace. The central problem is not only whether autonomous machines can perform agricultural tasks accurately. It is whether farmers, workers, contractors, and rural communities can understand, supervise, question, override, and hold these systems accountable when they act in complex farm environments. Without deliberate attention to trust calibration, human override, explainability, and accountability, autonomous farming may produce technically impressive but socially fragile systems. This article develops a Human-Centred Governance Framework for autonomous farming. The framework is designed around four pillars: calibrated trust, meaningful human override, explainable decision-making, and clear accountability structures. Its purpose is to shift governance attention from machine capability alone toward the conditions under which autonomous systems can be responsibly integrated into agricultural work. The proposed framework shows that human-centred governance is not an optional ethical supplement to autonomous farming. It is a practical requirement for safe deployment, legitimate adoption, and accountable operation. By integrating trust, override, explainability, and accountability, the framework provides guidance for farm managers, technology developers, regulators, and agricultural service providers.

Keywords Explainability, Autonomous farming, Agricultural robotics, Human-centred governance, Trust, Human override

*Correspondence:

Mina Takeda

takeda.m@kais.kyoto-u.ac.jp

¹ Department of Smart Agricultural Engineering, Graduate School of Agriculture, Kyoto University, Kyoto, Japan

Introduction

Autonomous farming is increasingly presented as a solution to agricultural labour shortages, productivity pressures, and the need for more precise management of crops, inputs, machinery, and field operations. Robotic harvesting, autonomous tractors, drone-based spraying, robotic weeding, and AI-supported sensing systems are part of a wider shift from tool-based mechanisation toward machine autonomy in agricultural work [1]. This shift raises expectations of efficiency and resilience, but it also changes the relationship between human judgement and machine action on farms.

The governance question created by autonomous farming is therefore not limited to whether robots can detect weeds, navigate rows, avoid obstacles, or execute tasks accurately. It also concerns who remains in control, who understands

the basis of machine decisions, and who is responsible when autonomous action causes crop damage, worker injury, animal disturbance, or environmental harm [2]. As agricultural robotics becomes more commercially viable, governance must address the conditions under which autonomy can be trusted, supervised, interrupted, and reviewed.

Existing discussions of robotics in agriculture show that adoption is shaped not only by technical performance but also by labour relations, farm routines, operator confidence, cost, social acceptance, and institutional readiness [3]. Human–robot interaction research further suggests that agricultural environments create distinctive challenges because machines operate in outdoor, variable, and often shared spaces rather than controlled industrial settings [4]. These challenges require governance frameworks that treat farmers and workers as active supervisors rather than passive recipients of automated decisions.

This article proposes a Human-Centred Governance Framework for autonomous farming systems. It integrates insights from agricultural robotics, responsible AI, safety engineering, meaningful human control, explainability, and trust in automation [5]. The article first frames autonomous farming as a governance challenge, then develops the pillars of trust, override, explainability, and accountability, before presenting an integrated framework and implementation recommendations.

Autonomous Farming as a Governance Challenge

Autonomous farming differs from conventional agricultural mechanisation because machines are no longer merely extending human physical capacity. They increasingly perceive farm conditions, classify objects, make operational choices, and execute actions with limited direct human involvement [6]. This creates a governance challenge because decision authority becomes distributed across farmers, software systems, robotic platforms, manufacturers, data infrastructures, and service providers.

Farm environments make this challenge especially difficult because they are open, changing, and biologically complex. Robots may encounter uneven terrain, variable light, mud, dust, livestock, workers, bystanders, machinery, weather shifts, and unpredictable crop conditions [4]. As a result, safety cannot be reduced to mechanical reliability alone, because responsible operation depends on how machines interpret context and how humans remain able to supervise those interpretations.

Figure 1 illustrates how autonomous farming shifts decision authority from conventional human-operated machinery toward distributed human–robot governance requiring trust calibration, override, explainability, and accountability.




Figure 1. From Machine Capability to Human-Centred Governance in Autonomous Farming Systems

The governance problem is also intensified by the pace at which commercial enthusiasm can outrun institutional preparedness. Studies of agricultural robotics and automation note that adoption barriers include cost, technical uncertainty, skill requirements, infrastructure gaps, and unclear responsibility structures [7]. These barriers are not merely market obstacles; they indicate that autonomous farming requires organizational, legal, and ethical arrangements capable of guiding deployment before accidents or legitimacy failures occur.

A further concern is that technology-centred safety standards may not fully address the human-centred dimensions of farm autonomy. Responsible robotics research argues that design guidance should consider the farming system as a whole, including operators, work practices, animal welfare, rural livelihoods, and downstream consequences [8]. Therefore, autonomous farming should be governed as a socio-technical system in which safety, trust, explainability, and accountability are designed together rather than treated as separate afterthoughts.

Trust in Agricultural Robotics

Trust is a central governance issue because autonomous machines must be relied upon without being accepted uncritically. In agricultural robotics, trust concerns whether farmers and workers believe that a system will perform reliably, behave predictably, and remain controllable under real farm conditions [9]. Trust is therefore not a vague attitude toward technology but a practical judgement about whether machine autonomy can be safely incorporated into daily agricultural work.

The goal of governance should be calibrated trust rather than maximum trust. Over-trust can encourage operators to ignore warning signs, delegate too much authority, or assume that the robot is safer and more competent than it actually is [10]. Under-trust can also be harmful because farmers may refuse useful automation, disable safety functions, or avoid systems that could support more sustainable and efficient operations.

Trust develops through repeated experience, visible reliability, understandable behaviour, and alignment with farm priorities. Human–robot collaboration studies show that operators need systems that communicate status, limits, uncertainty, and intended actions in ways that fit the tempo and practical realities of farm work [11]. Trust is also shaped by broader expectations about whether robotics supports farmers' expertise or displaces judgement, labour, and local knowledge [12].

A human-centred governance framework must therefore define trust as an outcome of design, training, communication, and accountability rather than as a burden placed on farmers alone. Ethical analyses of agricultural AI and robotics show that trust is weakened when systems are opaque, socially disruptive, or poorly aligned with farmer values and responsibilities [13]. **Table 1** outlines the dimensions of trust in agricultural robotics and the factors that build or erode it.

Table 1. Dimensions of Trust in Agricultural Robotics: Antecedents, Calibration Risks, and Implications for Human-Robot Collaboration

Trust dimension	Meaning in autonomous farming	Factors that build trust	Factors that erode trust	Governance implication
Reliability trust	Confidence that the robot performs the intended task consistently under farm conditions	Stable task performance, robust navigation, accurate sensing, predictable responses to field variability	Frequent faults, unexplained stops, inconsistent classification, poor performance in mud, dust, rain, or uneven terrain	Manufacturers and operators should document operational limits and validate performance across realistic farm scenarios

Predictability trust	Confidence that farmers can anticipate what the system is likely to do next	Clear machine status, visible operating mode, advance warnings, understandable movement patterns	Sudden movement, unclear mode changes, unexpected route choices, silent autonomy	Interfaces should communicate intention, operating mode, uncertainty, and next action before high-risk movement or intervention
Competence trust	Belief that the robot has the technical capability to perform the delegated agricultural task	Demonstrated accuracy, task-specific training, appropriate sensors, field-tested autonomy	Overgeneralised claims, poor weed or obstacle recognition, weak adaptation to local conditions	Deployment should match autonomy level to task complexity, field conditions, and operator skill
Control trust	Confidence that humans can intervene, pause, override, or halt the machine when necessary	Accessible emergency stop, clear override hierarchy, reliable remote supervision, fail-safe modes	Delayed stop response, hidden automation logic, unclear responsibility for intervention	Trust requires meaningful human control, not merely nominal supervision
Transparency trust	Confidence that the system can explain relevant actions and decisions	Farmer-facing explanations, logs, alerts, traceable decision histories	Black-box spraying, unexplained avoidance behaviour, inaccessible data records	Explainability should be built into routine operation, incident review, and training
Institutional trust	Confidence that manufacturers, regulators, insurers, and farm managers will act responsibly	Clear accountability, maintenance obligations, safety certification, farmer support	Liability ambiguity, weak regulation, poor after-sales support, unclear data ownership	Governance must allocate responsibility before deployment and provide review mechanisms after incidents
Social trust	Confidence that robotics fits local labour, community, and environmental expectations	Worker consultation, responsible innovation, gradual deployment, visible safety practices	Labour displacement anxiety, exclusion of farm workers, perceived loss of farmer agency	Adoption strategies should include social dialogue, training, and participatory design

Human Override and Operational Safety

Human override is the second pillar of the proposed framework because autonomy in farming should never remove the practical ability of humans to intervene. Meaningful human control requires more than placing a nominal operator near a machine; it requires that the operator can understand the situation, identify risk, and stop or redirect the system in time [14]. In autonomous farming, this principle is especially important because robots operate near crops, workers, animals, vehicles, buildings, waterways, and neighbouring land.

Operational safety must therefore be designed around realistic farm conditions rather than idealised assumptions about supervision. A farmer may be monitoring several machines, responding to livestock, managing irrigation, communicating with contractors, or moving between fields while an autonomous system is operating [15]. If override mechanisms are slow, hidden, unreliable, or difficult to access during mobile work, the system cannot be considered meaningfully controllable.

The safety challenge is also psychological because automation can change how operators pay attention. When machines perform well for long periods, farmers and workers may become less vigilant, especially if interfaces do not communicate uncertainty, boundary conditions, or emerging risk [16]. This creates the danger of automation complacency, where the human remains formally responsible but is practically unable to detect and correct failure at the moment when intervention is needed.

Human override also has legal and organizational dimensions. If a manufacturer claims that a machine is autonomous but expects the farmer to prevent all harmful outcomes, responsibility is displaced without being clearly governed [17]. A human-centred framework should therefore require explicit override rights, defined intervention procedures, training obligations, emergency-stop reliability, and documented responsibility for decisions made before, during, and after autonomous operation.

Explainability and Decision Transparency

Explainability is the third pillar because farmers need to understand not only what an autonomous system did but why it acted in a particular way. A robotic sprayer that treats one patch and avoids another, a weeding robot that removes some plants and preserves others, or a drone that changes its flight path all produce decisions that affect agronomic outcomes [18]. Without farmer-facing explanations, these systems risk becoming opaque authorities in environments where local judgement remains essential.

Explainable artificial intelligence research shows that transparency is not a single technical property but a relationship between model behaviour, user needs, decision context, and accountability requirements [19]. In autonomous farming, explanations must be practical rather than abstract. A farmer does not need a full mathematical reconstruction of a model; the farmer needs to know whether the system acted because of crop stress, weed classification, obstacle avoidance, sensor uncertainty, weather conditions, or safety limits.

Decision transparency also supports trust calibration. If a robot can indicate that it is uncertain about plant classification, operating near a boundary, or deviating from its planned path, the farmer can decide whether to continue, supervise more closely, or intervene [20]. This turns explainability into an operational safety function rather than a merely informational feature.

However, explainability has limits that governance must acknowledge. Post-hoc explanations can create a false sense of understanding if they simplify complex machine behaviour or present confidence where the system was actually uncertain [21]. For this reason, autonomous farming governance should require explanation logs, uncertainty communication, incident-review records, and design processes that align explanations with the actual decisions farmers must make in real time.

Accountability in Robotic Farm Systems

Accountability is the fourth pillar because autonomous farming distributes responsibility across several actors. When a robot damages crops, injures a worker, misapplies chemicals, frightens livestock, or causes environmental harm, it may be unclear whether responsibility lies with the farmer, manufacturer, software developer, maintenance provider, data supplier, or regulator [22]. This ambiguity can weaken safety because each actor may assume that another party is responsible for risk control.

Legal scholarship on small autonomous agricultural robots shows that existing liability frameworks may struggle to address machines that learn, adapt, or act without direct human command [17]. Traditional assumptions about operator control become less stable when autonomous systems make navigation, classification, and action decisions. Governance

must therefore define responsibility before deployment rather than waiting for courts, insurers, or regulators to interpret responsibility after harm occurs.

Accountability also concerns the internal organization of farming operations. Farm managers need clear rules for who authorizes autonomous operation, who monitors the machine, who responds to alerts, who documents incidents, and who decides whether a system can continue working after a near miss [23]. Without such rules, accountability remains informal, and safety becomes dependent on individual improvisation rather than governed practice.

Responsible innovation literature emphasizes that accountability should include social and ethical consequences, not only technical failure. Autonomous farming can affect labour roles, rural skills, data control, environmental practices, and farmer autonomy [24]. A human-centred governance framework should therefore connect incident liability with broader accountability for design choices, deployment conditions, training quality, maintenance, transparency, and stakeholder participation.

Proposed Human-Centred Governance Framework

The proposed Human-Centred Governance Framework integrates calibrated trust, meaningful human override, explainable decision-making, and clear accountability structures into a single governance model. These pillars are mutually reinforcing rather than separate safeguards. Trust enables effective collaboration, override provides a safety backstop, explainability supports both trust and intervention, and accountability creates incentives for safer design and more responsible deployment [1].

The first principle of the framework is that autonomous farming systems should be designed for supervised autonomy rather than unrestricted machine independence. This means that farmers and workers remain informed participants in autonomous operation, even when the machine performs tasks without continuous manual control [25]. Supervised autonomy does not reject innovation; it defines the conditions under which innovation remains compatible with human agency, safety, and responsibility.

Figure 2 presents the proposed Human-Centred Governance Framework as an integrated operational model linking farmers, manufacturers, regulators, and robotic systems through trust, override, explainability, safety, and accountability mechanisms.




Figure 2. Integrated Human-Centred Governance Framework for Trustworthy, Explainable, Overrideable, and Accountable Autonomous Farming

The second principle is role clarity across the lifecycle of the robotic system. Farmers should be responsible for appropriate use, site-specific risk assessment, and trained supervision, while manufacturers should be responsible for

safety design, reliable override mechanisms, transparent documentation, and post-deployment support [26]. Regulators and industry bodies should define baseline safety expectations, certification pathways, incident-reporting requirements, and minimum standards for explainability and human control.

The third principle is that governance should be operationalized through practical checklists, training routines, logs, and review mechanisms rather than expressed only as abstract ethical commitments. Agricultural robots should be assessed before deployment, monitored during operation, and reviewed after incidents or near misses [27]. **Table 2** presents the proposed Human-Centred Governance Framework integrating the four pillars.

Table 2. Proposed Human-Centred Governance Framework for Autonomous Farming: Pillars, Principles, and Operational Requirements

Governance pillar	Core governance principle	Farmer and operator requirements	Manufacturer and developer requirements	Regulator and industry requirements	Operational evidence required
Calibrated trust	Trust should match the system’s demonstrated capability, limitations, and operating conditions	Use the system only within validated conditions, monitor uncertainty warnings, report reliability concerns, avoid blind reliance	Communicate performance limits, provide transparent operating modes, disclose failure conditions, design farmer-facing status displays	Require realistic field validation, prohibit exaggerated autonomy claims, support trust-related safety guidance	Field validation records, operator training logs, reliability reports, uncertainty alerts, user feedback records
Meaningful human override	Humans must retain practical authority to pause, redirect, or stop autonomous operation	Know override procedures, maintain access to emergency stop functions, supervise high-risk operations, document interventions	Provide accessible stop controls, fail-safe modes, remote intervention capacity, clear override hierarchy, rapid response design	Set minimum override standards, certify emergency-stop performance, require human-control documentation	Override test records, emergency-stop response times, intervention logs, incident simulations, maintenance checks
Explainable decision-making	Farmers should be able to understand relevant machine actions, classifications, and safety decisions	Review explanations for critical actions, use explanation logs in operational decisions, challenge unexplained behaviour	Provide interpretable alerts, decision logs, uncertainty communication, farmer-oriented explanations, accessible documentation	Define minimum explanation requirements for high-risk autonomous farming tasks	Explanation logs, decision traces, classification confidence records, operator review notes, post-incident explanation reports
Operational safety	Safety should be treated as a continuous socio-technical process, not a one-time technical feature	Conduct site risk assessments, manage human and animal proximity, supervise	Design for outdoor variability, provide safety cases, update systems responsibly, maintain hardware	Require safety certification, incident reporting, near-miss learning systems, and	Site risk assessments, pre-operation checklists, maintenance records, near-miss

		hazardous tasks, stop operation under unsafe conditions	and software reliability	cross-sector safety standards	reports, safety audit outcomes
Clear accountability	Responsibility must be allocated before deployment and traceable after harm or near misses	Define who authorizes use, who monitors, who intervenes, and who reports incidents	Clarify warranty, update responsibility, software limits, data access, and support obligations	Clarify liability expectations, insurance requirements, audit duties, and reporting obligations	Responsibility matrix, service agreements, insurance documentation, incident reports, audit trails
Participatory governance	Farmers, workers, and affected stakeholders should influence deployment conditions	Consult workers and contractors, incorporate local knowledge, identify community concerns, adapt deployment plans	Involve users in design testing, usability evaluation, and safety feedback processes	Encourage participatory standards development and responsible innovation platforms	Consultation records, design feedback reports, training attendance, stakeholder review documentation

Implementation Risks and Managerial Recommendations

Implementation will be difficult because human-centred governance creates costs that may not be immediately visible in productivity calculations. Explainability interfaces, operator training, safety validation, documentation systems, and incident review processes require investment beyond the purchase price of the machine [28]. Farm managers should therefore evaluate autonomous systems through total governance cost, including training, supervision, maintenance, insurance, and accountability infrastructure.

A second risk is uneven readiness across farms, manufacturers, and regulatory systems. Smaller farms may lack the technical staff to evaluate autonomous systems, while manufacturers may vary in the quality of their documentation, support, and safety design [7]. Managers should begin with supervised autonomy, limited deployment zones, staged operational testing, and clear rules for when machines must be paused, manually supervised, or removed from service.

A third risk is regulatory fragmentation, especially when autonomous machinery, AI decision-making, chemical application, occupational safety, environmental protection, and data governance are handled through separate institutional channels. Responsible AI in agriculture requires systemic risk understanding because technical decisions can create labour, environmental, legal, and social effects beyond the immediate task [5]. For this reason, the sector needs dialogue platforms connecting farmers, manufacturers, regulators, insurers, researchers, and rural communities around shared standards for trust, override, explainability, and accountability.

Conclusion

Autonomous farming can contribute to more efficient, precise, and resilient agricultural systems, but only if its deployment remains centred on the humans who work with, supervise, and are affected by robotic machinery. The central governance challenge is not simply how to make agricultural robots more capable. It is how to ensure that capability remains safe, understandable, controllable, and accountable in real farm environments.

This article proposed a Human-Centred Governance Framework built on four pillars: calibrated trust, meaningful human override, explainable decision-making, and clear accountability structures. Together, these pillars shift autonomous

farming governance away from a narrow focus on machine performance and toward a broader concern with human agency, operational safety, social legitimacy, and responsible innovation.

The framework provides a practical roadmap for farm managers, manufacturers, regulators, and industry bodies. Its adoption would help ensure that autonomous farming systems are not merely efficient but also trustworthy, transparent, interruptible, and accountable. Human-centred governance should therefore be treated as a foundational condition for the responsible future of robotic agricultural systems.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 23 Dec 2023

Revised: 06 Feb 2024

Accepted: 10 Mar 2024

Published online: 15 September 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Rose DC, Lyon J, de Boon A, Hanheide M, Pearson S. Responsible development of autonomous robotics in agriculture. *Nat Food*. 2021;2(5):306-9.
<https://doi.org/10.1038/s43016-021-00287-9>.
- Sparrow R, Howard M. Robots in agriculture: Prospects, impacts, ethics, and policy. *Precis Agric*. 2021;22(3):818-33.
<https://doi.org/10.1007/s11119-020-09757-9>.
- Legun K, Burch K. Robot-ready: How apple producers are assembling in anticipation of new AI robotics. *J Rural Stud*. 2021;82:380-90.
<https://doi.org/10.1016/j.jrurstud.2021.01.032>.

Vásconez JP, Kantor GA, Auat Cheein FA. Human–robot interaction in agriculture: A survey and current challenges. *Biosyst Eng*. 2019;179:35–48.

<https://doi.org/10.1016/j.biosystemseng.2018.12.005>.

Tzachor A, Devare M, King B, Avin S, Ó hÉigeartaigh S. Responsible artificial intelligence in agriculture requires systemic understanding of risks and externalities. *Nat Mach Intell*. 2022;4(2):104–9.

<https://doi.org/10.1038/s42256-022-00440-4>.

Yang Q, Du X, Wang Z, Meng Z, Ma Z, Zhang Q. A review of core agricultural robot technologies for crop productions. *Comput Electron Agric*. 2023;206:107701.

<https://doi.org/10.1016/j.compag.2023.107701>.

Gil G, Casagrande DE, Cortés LP, Verschae R. Why the low adoption of robotics in the farms? Challenges for the establishment of commercial agricultural robots. *Smart Agric Technol*. 2023;3:100069.

<https://doi.org/10.1016/j.atech.2022.100069>.

Eastwood CR, Dela Rue B, Edwards JP, Jago J. Responsible robotics design—A systems approach to developing design guides for robotics in pasture-grazed dairy farming. *Front Robot AI*. 2022;9:914850.

<https://doi.org/10.3389/frobt.2022.914850>.

Benos L, Moysiadis V, Kateris D, Tagarakis AC, Busato P, Pearson S, et al. Human–robot interaction in agriculture: A systematic review. *Sensors (Basel)*. 2023;23(15):6776.

<https://doi.org/10.3390/s23156776>.

Kohn SC, de Visser EJ, Wiese E, Lee YC, Shaw TH. Measurement of trust in automation: A narrative review and reference guide. *Front Psychol*. 2021;12:604977.

<https://doi.org/10.3389/fpsyg.2021.604977>.

Yerebakan MO, Hu B. Human–robot collaboration in modern agriculture: A review of the current research landscape. *Adv Intell Syst*. 2024;6(7):2300823.

<https://doi.org/10.1002/aisy.202300823>.

Marinoudi V, Sørensen CG, Pearson S, Bochtis D. Robotics and labour in agriculture: A context consideration. *Biosyst Eng*. 2019;184:111–21.

<https://doi.org/10.1016/j.biosystemseng.2019.06.013>.

Dara R, Hazrati Fard SM, Kaur J. Recommendations for ethical and responsible use of artificial intelligence in digital agriculture. *Front Artif Intell*. 2022;5:884192.

<https://doi.org/10.3389/frai.2022.884192>.

Santoni de Sio F, van den Hoven J. Meaningful human control over autonomous systems: A philosophical account. *Front Robot AI*. 2018;5:15.

<https://doi.org/10.3389/frobt.2018.00015>.

Lytridis C, Kaburlasos VG, Pachidis T, Manios M, Vrochidou E, Kalampokas T, et al. An overview of cooperative robotics in agriculture. *Agronomy*. 2021;11(9):1818.

<https://doi.org/10.3390/agronomy11091818>.

Shneiderman B. Human-centered artificial intelligence: Reliable, safe and trustworthy. *Int J Hum Comput Interact*. 2020;36(6):495–504.

<https://doi.org/10.1080/10447318.2020.1741118>.

Basu S, Omotubora A, Beeson M, Fox C. Legal framework for small autonomous agricultural robots. *AI Soc*. 2020;35(1):113–34.

<https://doi.org/10.1007/s00146-018-0846-4>.

Arrieta AB, Díaz-Rodríguez N, Del Ser J, Benetot A, Tabik S, Barbado A, et al. Explainable artificial intelligence: Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion*. 2020;58:82-115.
<https://doi.org/10.1016/j.inffus.2019.12.012>.

Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F, Pedreschi D. A survey of methods for explaining black box models. *ACM Comput Surv*. 2018;51(5):93.
<https://doi.org/10.1145/3236009>.

Sparrow R, Howard M, Degeling C. Managing the risks of artificial intelligence in agriculture. *NJAS Impact Agric Life Sci*. 2021;93(1):172-96.
<https://doi.org/10.1080/27685241.2021.2008777>.

Ryan M. The social and ethical impacts of artificial intelligence in agriculture: Mapping the agricultural AI literature. *AI Soc*. 2023;38(6):2473-85.
<https://doi.org/10.1007/s00146-021-01377-9>.

Ryan M, van der Burg S, Bogaardt MJ. Identifying key ethical debates for autonomous robots in agri-food: A research agenda. *AI Ethics*. 2022;2(3):493-507.
<https://doi.org/10.1007/s43681-021-00104-w>.

van der Burg S, Bogaardt MJ, Wolfert S. Ethics of smart farming: Current questions and directions for responsible innovation towards the future. *NJAS Wagening J Life Sci*. 2019;90-91:100289.
<https://doi.org/10.1016/j.njas.2019.01.001>.

Bronson K. Looking through a responsible innovation lens at uneven engagements with digital farming. *NJAS Wagening J Life Sci*. 2019;90-91:100294.
<https://doi.org/10.1016/j.njas.2019.03.001>.

van der Burg S, Giesbers E, Bogaardt MJ, Ouweltjes W, Lokhorst K. Ethical aspects of AI robots for agri-food; a relational approach based on four case studies. *AI Soc*. 2024;39(2):541-55.
<https://doi.org/10.1007/s00146-022-01429-8>.

Burton S, Habli I, Lawton T, McDermid J, Morgan P, Porter Z. Mind the gaps: Assuring the safety of autonomous systems from an engineering, ethical, and legal perspective. *Artif Intell*. 2020;279:103201.
<https://doi.org/10.1016/j.artint.2019.103201>.

Ditzler L, Driessen C. Automating agroecology: How to design a farming robot without a monocultural mindset? *J Agric Environ Ethics*. 2022;35(1):2.
<https://doi.org/10.1007/s10806-021-09876-x>.

Lowenberg-DeBoer J, Huang IY, Grigoriadis V, Blackmore S. Economics of robots and automation in field crop production. *Precis Agric*. 2020;21(2):278-99.
<https://doi.org/10.1007/s11119-019-09667-5>.