

---

## **AI Ethics in Autonomous Agricultural Machinery**

**Dr. Adrian M. Kellerman**

Associate Professor, Department of Agricultural Robotics and Intelligent Systems  
Great Plains Institute of Technology, Lincoln, Nebraska, USA

**Corresponding Author: Dr. Adrian M. Kellerman**

Associate Professor, Department of Agricultural Robotics and Intelligent Systems  
Great Plains Institute of Technology, Lincoln, Nebraska, USA

---

**Article Type: Simulation-Based Research Article**

**Submitted:** 04-February-2025 **Revised:** 11-February-2025

**Accepted:** 18-February-2025 **Published:** 25-February-2025

---

**Abstract:** Artificial intelligence is gradually changing agricultural motorized equipment from simple operator-controlled machinery into more autonomous systems capable of autonomous perception, navigation, obstacle avoidance, precision spraying, mechanized weeding, crop surveillance, harvesting assistance and fleet level decision support. Autonomous agricultural machinery has the potential to greatly increase operational efficiency, improve labor shortages, reduce unnecessary chemical use and allow for more precise farming practices. However, increasing machinery autonomy can pose ethical and governance challenges that cannot be solved with simple performance metrics.

Agricultural robots operate in inconsistent, uncertain, constantly changing, safety critical environments with farm workers, livestock, crops, infrastructure, airborne dust, variable illumination, uneven terrain and unreliable network connectivity. Failures in autonomous decision making or perception can create physical, environmental, economic, cyber security and accountability risk. Here I develop a simulationgenerated Responsible-AI framework for autonomous agricultural machinery that co-optimizes autonomous perception for safety and uncertainty; redundancy and recall; human override and fail-safe design; environmental stewardship; interpretability and auditability; and fail-safe operation. A synthetic audit consisting of 2,400 autonomous field-operation episodes was created and analyzed across four autonomous operation governance architectures Automation-First; Safety-Constrained AI; Human-in-the-Loop; and Responsible AI. No real farm accidents, worker injury reports, manufacturer datasets, agribusiness telemetry or farm market data were used.

The Responsible-AI configuration achieved a 0.97 obstacle-detection recall, a human-override success rate of 97.3%, with environmental-intervention precision of 95.2/100 and modeled high-risk-event rate of 2.6%, compared to 12.0% under Automation-First operation. Ordinal logistic regression identified presence of workers or animals, sensor discrepancy and low visibility as strong positive predictors of high-risk-events, whereas mandated fail-safe stopping, human override, redundant perception, and dynamic, environment-sensitive speed reduction demonstrated large protective effects. These results suggest that ethical autonomy in agriculture should not seek maximum independent machine operation. Instead, responsible agricultural AI should employ experience-driven, situation-aware autonomous operation in which operational authority is dynamically adjusted based on environmental uncertainty, safety significance, sensor reliability and the availability of appropriate human oversight.

**Keywords:** autonomous farming machines, responsible AI, agricultural robots, AI ethics, precision farming, AI human supervision, intelligent machines, safe computing machines, smart farms, AI explanation.

## Introduction

The evolution of agricultural machinery has been a long-term trend toward digital integration of machine control and the use of satellite navigation, machine vision, automatic steering, sensor fusion, artificial intelligence, variable-rate application, and independence in control systems. Research in agricultural robotics indicated support for autonomous navigation, field operations, monitoring, spraying, mechanical treatment and other advanced precision applications, and autonomous agricultural vehicles include perception, localization, planning, and actuation in complex autonomous control architectures.

Autonomy is an appealing feature for various reasons. Farms are under pressure to improve efficiency, while simultaneously minimizing excess consumption of resources, and increasing the accuracy of applications, in conjunction with labor constraints and the demands of more complex production systems. Agricultural robots are thus an integral hardware feature of digital farming and Shamshiri and co-authors frame agricultural robotics as being a part of a larger digital transition in agriculture, while Bechar and Vigneault point to the challenge of the incorporation of robotic subsystems into robust field-platforms.

However, autonomous agriculture is more profoundly distinct from automation seen in closely managed industrial settings. Farming machinery must operate successfully on inconsistent soil, crop rows, slopes, dust, mud, variable lighting, greenery, livestock, human laborers, and temporarily displaced objects. Perception apparatus that is robust in one operational context may be inadequate at best when occlusion, poor observing conditions, incompatible granule, or unpredicted alteration is encountered. Studies of multisensor navigation exemplify the reason that agricultural robots frequently use multiple combinations of vision, LiDAR, ultrasonic, and other power sources.

The ethical issues, therefore, are not just whether autonomous machines can each fulfill their tasks but whether a purportedly excellent machine is ethically unacceptable if it cannot differentiate people reliably, if any uncertainty is hidden from the machine's operators, if the emergency stop cannot kill the machine, if its field data are transferred into the wrong hands, or if responsibility is forsaken once the machine get automated.

Several ethical problems are particularly important:

- **Safety:** autonomous machinery may cause the following hazards; collision, crushing, entanglement or hazards from an implement.
- The **management** of uncertainties: perception performance often degrades in dust, dark, occlusion crop or disagreement of sensors.
- **Human** SOU: operators should have a substantial degree of handson control over high-consequence actions of the machine.
- **Accountability:** At least for these people involved; machine producer, software delivery, owners/operators to be accountable.
- Protection of the environment: off-target effect can be occurring due to perception or localizations in case of automated spraying, fertilization and mechanical control.
- **Data governance:** interconnected equipment may share commercially sensitive farm, yield, field and operational data.
- **Cyber security:** remote access can provide potential routes for machine operation and/or farm data to be corrupted.
- **Labor implications:** Robotics in agriculture could impact job content, skill demands, and reliance on technology providers.

The motivation for this study is thus to progress from the ‘should an agricultural machine be autonomous?’ question to the more meaningful ‘when should autonomy be allowed to persist?’ question.

The 2018 version of ISO 18497 offered design recommendations for highly automated agricultural farm machinery, and explored safe working practices and residual risk. This showed that before the advent of autonomous operating agriculture, any safety considerations were to be viewed as requiring their own focus independent of general machinery. This thesis now takes this reasoning further in designing a Responsible-Ai model where safety and ethical governance are early-elevated autonomous authority.

## Literature Review

### a) Autonomous Agricultural Machinery and Robotic Field Operations

Agricultural robotics is a broad field, with many different robotic platforms and functions. A seminal review by Bechar and Vigneault categorizes agricultural robots as a set of integrated components: sensors, decision modules, actuators, and mobile bases. This later systems-based review highlighted that operation and whole-system integration take precedence over the focus on individual sensing or control modules:

Likewise, autonomous agricultural vehicles have benefited from advancements in navigation, localization, path planning, obstacle detection, and control techniques using machine vision and satellite navigation. A 2020 review by Roshanianfard and Noguchi provides a comprehensive summary of the advancements in autonomous agricultural vehicles, underscoring the ongoing need for dependable navigation and field-specific system design. Research in computer-vision has further broadened the scope of automation to encompass tasks such as crop detection, system navigation, quality inspection, disease detection, and robotic manipulation.

In this literature it appears that the Autonomous performing a Machine is through reliance on complex multiple-dimensioned technical competencies:

- robust localization
- Senses of people, animals, crops, and obstacles
- Created the reliable sensor fusion and
- Safe path planning

- task-specific actuation
- machine-state monitoring
- communication integrity
- fault detection
- emergency stopping
- The overall recovery: the operator is primarily concerned with recovering to the scheduled operation after a failure.

Apart from the individual component failure, the malfunction can spread to the entire machine system. For instance, an autonomous tractor may navigate properly but unsafely at the same time if the obstacle perception is compromised; An intelligent sprayer may precisely recognize weed patches but cause environmental damage if the localization drift lead to spraying off target.

## **b) Safety, Human–Robot Interaction, and Operational Uncertainty**

Safety is especially difficult in agricultural robotics because the environment cannot be fully designed for the machine. The 2018 ISO 18497 standard established design principles for highly-automated agricultural machinery and covered the importance of accounting for significant hazards, expected operation and safe working practices. Work reviewing safety requirements for autonomous machines has also focused on obstacle detection, control-system validity, safety-related functionality and the ongoing challenge of establishing assurance in real agricultural environments.

Interaction between humans and robots is just as significant. Human workers will be in close proximity with the harvesting robot systems, tractors, carriers, as well as semi- or fully-autonomous robots. The work done by Benos et al. on safety and ergonomics of human–robot interaction in agriculture. illustrates the notion that autonomy is simply a function of the technology and shows that one must always consider the human:

From existing literature, three principles can be derived. 1) the paradigm of farm safety should be context sensitive, since the operational risk varies with factors such as the speed, visibility, intimacy, terrain, task, implement, etc.; 2) autonomous systems needs to be fault tolerant, because individual perceptual mechanisms cannot be assumed to be trustworthy in all environments, and 3) the human supervision should not be viewed only as an emergency

backup in the later stage of system development, but the basis for design and limitation of autonomy.

These findings support an adaptive-autonomy approach. The same machine in an open environment, with robust redundant perception with high confidence should be entrusted with operational authority. If the same machine enters an environment with poor visibility, or populated with human workers, the machine should slow, increase safety zones, increase attentiveness or hand the controls to a human.

### **c) Smart-Farming Ethics, Data Governance, and Responsible Innovation**

Beyond physical safety concerns, similar ethical factors may be relevant in AGI, including data ownership and availability, and horizontal and vertical distribution of power and benefits. According to Van der Burg, Bogaardt, and Wolfert, these are thus far identified as significant themes in smart-farming ethics. As Wolfs et al's show, big-data agriculture thus generates both technical opportunities and socioeconomic issues with regard to data platforms, governance arrangements, and stakeholder relations.

Jakku et al. show that trust towards smart farming can be built by informing farmers on data use and benefit sharing. Klerkx et al. et al. show that digital agriculture bears social implications in different domains such as knowledge, skills, farm management, power and changing institutional arrangements. These issues can be transposed to autonomous machinery, a modern agricultural robot can be in the same time a machine, a platform of sensors, a decision system and the infrastructure of data generation.

Responsibility becomes a legal challenge too. Basu et al explored the legal framework for small autonomous agricultural robots and demonstrated how regulation functions as a barrier or facilitator, contingent on the regulation of autonomous machines. Relf-Eckstein et al also looked to autonomous farm machinery in the broader context of innovation and policy, showing that innovation itself is likely to outrun the institutional context.

### **Research Gap**

Many concepts from the field of agricultural-robotics literature can be utilized in terms of navigation, perception, mobile platforms, task execution, machine vision, and automation in the framing of hybrid systems. However, while the technical capability and ethical

governability are studied separately, the mind set of having such a gap between the autonomous systems and systems that the autonomous governing is properly limited has been manifested.

The first shortfall is the preponderance of task-performance comparisons. Agricultural robots tend to be tested by way of navigation accuracy, operational speed, object detection capability, task effectiveness, field capacity or application accuracy. Those are essential engineering metrics, but they cannot stand by themselves to assess whether the machine will behave safely in the face of sensor confidence drops or changing conditions.

Another area of lack is uncertainty-sensitive autonomy. Image processing algorithms usually generate a level of confidence, but conceptually, machine authority always falls into two bins: manual or autonomous. A more responsible system architecture would enable autonomous operation in the presence of high confidence, and which would decrease machine authority when there is uncertainty.

The third gap refers to the usefulness of human oversight. Simply putting in an emergency-stop button will not generate a valid oversight. Operators must be given time, access to information, an understandable interface, and the opportunity to intervene before the automated mistake becomes a real injury.

The last gap relates to environmental ethics. Autonomous sprayers and treatment equipment could improve input efficiency through precise application, but failures may also cause chemical drift and crop injury, soil disruption, or dual microbe impacts. Hence the environmental performance indicator be better viewed as an ethic contribution, instead of simply an agronomic efficiency indicator.

Fifth gap: data and control governance. Autonomous connected machinery might send location, field, crop, yield, machinery-performance, and operational data. Digital farming research has highlighted issues of trust, data ownership, transparency, and power sharing as ongoing issues in farm digitalization.

The study aims to fill these gaps in the context of a framework where autonomous agricultural equipment is assessed in terms of their productivity, accuracy of perception, controllability by operators, environmental accuracy, uncertainty handling and governance risk.

## Novelty / Proposed Contribution

This study puts forward a Responsible Agricultural Autonomy Framework that is based on the premise that autonomous authority should be bounded, not free.

The first is risk-adaptive autonomy. Machine speed, implement-waiting, freedom of path planning, and confirmation-by-human are associated with sources of environmental uncertainty and their consequences.

Second, the framework deploys redundant perception as a governance mechanism. Sensor fusion is not just considered as a mean to enhance average detection performance. Disagreement among sensors equates to uncertainty which can curtail machine behavior.

Third, the framework decouples “Task” completion from “Safe” task completion. A machine will not be recognized for positive performance if it quickly completes its task with high near-miss or intervention risk.

Fourth, the model combines environment-intervention accuracy. Foreign-interventions in the field such as spraying, cutting, ploughing should also considered their side-effect to other thing besides focusing on gains.

Fifth, human override is set up as an active safety variable. The schema considers the presence, timeliness, comprehensibility and efficacy of intervention.

Sixth, auditability and accountability are ingrained in the operation of machine. System keep the record of sensor willingness, automated decision, override use, emergency stop, and huge abnormality.

The Responsible-AI architecture is therefore based on several core principles:

- Since the more choice-is-related the less autonomy should be, the more uncertainty of the choice options increased.
- The more critical the decision, the greater the evidence necessary to substantiate this decision;
- , a better rule of thumb for indicating the presence of sensor disagreement is ‘rule or beacon of caution, not mere averaging’ (blow flies).
- Of particular concern is the need to maintain human intervention on high risk operations.
- “Controlled stop rather than speculative continuation should be the default fail-safe behavior;”

- The levels are anything but responsible. The above level of environmental degradation; it is a very serious crime; levels are not evidence of individual but rather they reflect a failure of the whole system:
- data collection should remain proportionate to operational purpose
- .... should record information that are relevant and useful for such purpose of incident analysis and responsibility.

## Research Objectives / Research Questions

The main goal is to design and test a simulation-oriented system of ethical governance applied to autonomous agricultural equipment.

The precise goals are to compare autonomous-machinery architectures by their performance in such fields as efficiency, perception dependability, human controllability, environmental accuracy, and probability of dangerous events; to test the hypothesis that hard fields increase modeled safety risk; to test the hypothesis that fail-safe stopping and human supervision decrease modeled safety risk; and to determine how best to include ethical constraints without discarding useful levels of machine autonomy.

The study addresses the following questions:

**RQ1:** How does Responsible-AI governance change operational efficacy vis-à-vis automation-first machinery?

**RQ2:** Do redundant sensing, human oversight, and fail-safe behavior argue against the modeled high-risk agricultural incidents?

**RQ3:** Operating conditions that most increase safety risk.

**RQ4:** Is it possible to have responsible autonomy in the design and operation of an automatic obstacle detection/avoidance system such that the system's reliability and acceptance is increased, as well as its environmental- intervention accuracy?

**RQ5:** What are the most significant protective relations of governance controls in the high-risk-event simulation model?

## Methodology

### a) Research Design

A simulation based comparative audit design was used. Four architectures of Agricultural machines were tested: Automation-First, Safety-Constrained AI, Human-in-the-Loop and Responsible AI.

This experiment was synthetic. It didn't analyze any real field worker, autonomous tractor, equipment manufacturer, farm accident, animal was injure, chemical plant, field telemetry, proprietary control software, commercially equipment data sets.

The intention was methodological: assessing to what extent and in which direction the increase of strength of the control mechanisms would affect a specified number of simulated autonomous operations.

### b) dataset/sample Choosing

The correct sample and understanding the sample was a further important factor in improving accuracy. Although the method of selecting the sample was largely random, some factors such as choosing only samples from the beginning of each category or only using data samples from hotter days appeared to slightly decrease accuracy.

The analytical sample consisted of 2400 synthetic autonomous field-operation episodes, 600 episodes each for the four governance architectures.

Each architecture was evaluated equally across four operating conditions:

1. Open and dry fields;
2. Establish soil depth and nutrient conservation/control; describe the species in the nutritive cycle to protect a soil characterized by clumping;
3. Tolling, if set above 50dB(A), subject to inclusion of the existing background noise levels;
4. Worker or animal proximity.

Hence each architecture-condition cell consisted of 150 simulated episodes.

**Table 1. Operational Variables in the Autonomous Agricultural Machinery Audit**

| Variable                | Operational Definition                                                                                                          | Measurement                                                                                                                                                                                                                           | Analytical Function |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| Field Operation         | Episode of autonomous machine Synthesis                                                                                         | N = 2,400                                                                                                                                                                                                                             | Unit of analysis    |
| Operating Condition     | Four environmental/safety conditions:                                                                                           | Categorical                                                                                                                                                                                                                           | Domain factor       |
| Sensor Confidence       | Confidence for perceived objects                                                                                                | 0–1                                                                                                                                                                                                                                   | Autonomy control    |
| Sensor Disagreement     | Channel disagreement in perception                                                                                              | 0 / 1 As can be seen from the above; there appears to be instability in the system; the Bode plot shows that the gain will be increased to be closer to 0dB and the phase needs to be improved to get an approximation of 0o at 1kHz. | Risk trigger        |
| Worker/Animal Proximity | Location within defined safety region                                                                                           | 1 0 / 1. While . 1 0 / 1 might be better represented as but if we take just, then.                                                                                                                                                    | Safety variable     |
| Dynamic Speed Reduction | Automatic reduction in the presence of risk.                                                                                    | 0 / 1. Occasional (1990) has been so far the only success of the underground banding movement in Finland.                                                                                                                             | Governance control  |
| Human Override          | Operator intervention success. Problems occurred when inexperienced operators were called in to deal with an unusual situation. | 0 / 1 is an evaluation method, defined as a, “method of evaluation similar in design to the overall marking unit”. Where I / 1 are the number of marks originally attained in a                                                       | Governance control  |

|                         |                                                                                                                                                                                                                                                                                                                                 |                                                                                                                                                                                                                                               |                                       |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|
|                         |                                                                                                                                                                                                                                                                                                                                 | piece of work and 0 / 1 are the number of marks ultimately awarded to give a numerical representation of the quality of work for a piece.                                                                                                     |                                       |
| Fail-Safe Stop          | Controlled safety stop.2 An underwater controlled safety stop (CS3) is often conducted on deep wreck dives in conjunction with sound underwater navigation during the ascent. It is an extremely useful practice in terms of flexibility of bottom time, although it may not be officially sanctioned by all training agencies. | Can you” because of respect for personal choice in treatment. However, the decision to seek treatment can be influenced by various factors including the individual’s motivation, age, operational status, and the severity of their failure. | Safety control by a primary incident. |
| Environmental Precision | Accurate targeted field treatment, therefore, is required for the effective control of pasture weeds.                                                                                                                                                                                                                           | 0–100                                                                                                                                                                                                                                         | Sustainability outcome                |
| High-Risk Event         | Modeled unsafe or materially adverse event                                                                                                                                                                                                                                                                                      | 0 / 1 implies the value 0. This means top is acceptable in the sense that it evaluates to 0, although d and g are not acceptable since they evaluate to 1.                                                                                    | Ethical outcome                       |

Table 1 is the operationalization of the technical and ethical variables. Sensor confidence and disagreement operationalize the technical conditions of computer perception, while worker or animal proximity operationalizes the condition that can make a common navigation error safety critical. Dynamic speed reduction, human override, and fail-safe stopping are operationalized as forms of governance control, because they are decisions about what the machine is allowed to do after uncertainty has been registered. Environment quality operationalizes the evaluation beyond collision safety as a measure of whether field action is still accurately targeted.

### **c) Usage of tools/ techniques of implementation include**

The simulated autonomous system included conceptual multi-sensor perception, route planning, object detection, task-control logic and risk-based intervention. The methodology is consistent with the agricultural-robotics literature showing that effective field operation requires multi-sensing, navigation and control sub-systems.

The analysis included:

- descriptive comparative analysis
- obstacle-detection recall
- Field-task efficiency scoring
- human-override success
- high-risk-event frequency
- environmental-intervention precision
- A binary logistic regression
- odds-ratio estimation
- 95% confidence interval
- relative-risk reduction

Furthermore, the configuration of Responsible-AI included fictitious audit logging in order to reproduce the log of critical actions performed by the system after an event occurred.

### **d) Experimental Procedure**

Each synthetic episode started by inputting to the machine a specific field task and route. The simulated perception system created information about field boundaries, crop rows, obstacles, people or animals, as well as weather conditions.

During Automation-First operation, the system would remain autonomous unless it reached some emergency levels of threshold. Safety-Constrained AI provided less aggressive obstacle thresholds and speed reduction.

High-uncertainty/safety-critical events required a secondary human acknowledgement for Human-in-the-Loop operation. Responsible AI integrated human oversight with redundant sensing, a priori sensor-disagreement detection, explainable alerts, variable speed control, limiter-enabled implement activation, fail-safe stopping, and audit-logging after actions.

The decision logic followed a risk-sensitive hierarchy:

- **Very low risk:** represents “regular or normally autonomous”, immediate to long-term autonomous function.
- **Moderate risk** — lower speed but with increased sensor checks;
- Limited autonomy (warning placard); minimized operator authority; operator reacts to warning. High uncertainty if:
- **Safety-critical uncertainty:** controlled stop and confirmation by human;
- **Confirmed hazard** : fail safe activated.

In so doing, the architecture implicitly does not presuppose that maximized autonomy should always be the goal.

## Proposed Questionnaire for Future Empirical Validation

The next five questions proposed for future studies with agricultural-machine operators were not asked in this current simulation:

1. Did the autonomous machinery make it clear why it decreased speed, came to a stop, or asked the human to intervene?
2. Would you have been able to switch back to manual control quickly and reliably when you thought automated control would be unsafe?
3. Were you aware of the specific field and of the data that was collected by the machine during the autonomous operation?
4. Did the machine warn us appropriately when sensors or automated decisions were uncertain?
5. Would having accessible and clear event logs and human override make you more confident to use autonomous farm equipment?

No survey statistics. The survey questions are still proposed instruments for research.

## e) Evaluation Metrics

Field-task efficiency was rated on a simulated 0–100 scale and included as a function of time to complete, route continuity and task completion.

Obstacle-detection recall were the proportions of relevant hazards detected by the simulated perception system. More was better.

Human-override (1) success the percentage of initiated operator interventions that successfully resulted in the machine being placed safely under human control.

Environment-intervention accuracy was the closeness of the simulated treatment or mechanical operation to the correct field target other than off-target locations.

The high-risk-event rate was the main measure of the ethical-risk. It consisted of simulated episodes of min unsafe proximity, uncontrolled operation under severe uncertainty, off-target environmental intervention, unsuccessful override, or failure to achieve a safe state when needed.

## Results

The model showed us tradeoffs between Freedom to automate and governability to search. Automation-First machinery received the highest raw task-efficiency score, but this came with weaker obstacle detection, much lower successful human intervention, lower environmental precision and the highest modeled high-risk-event rate. Responsible AI had marginally lower raw efficiency and much better measures of the broader safety and governance portfolio.

**Table 2. Simulated Performance Across Autonomous Agricultural Machinery Architectures**

| Architecture          | Field-Task Efficiency /100 | Obstacle-Detection Recall | Human-Override Success | Environmental Precision /100 | High-Risk Event Rate |
|-----------------------|----------------------------|---------------------------|------------------------|------------------------------|----------------------|
| Automation-First      | 91.2                       | 0.88                      | 62.4%                  | 82.1                         | 12.0%                |
| Safety-Constrained AI | 89.7                       | 0.93                      | 78.6%                  | 88.3                         | 6.8%                 |

|                   |      |             |              |             |             |
|-------------------|------|-------------|--------------|-------------|-------------|
| Human-in-the-Loop | 86.4 | 0.95        | 94.1%        | 91.0        | 3.5%        |
| Responsible AI    | 88.6 | <b>0.97</b> | <b>97.3%</b> | <b>95.2</b> | <b>2.6%</b> |

Table 2 shows the automation first machinery has a field-task efficacy of 91.2, compared to 88.6 under ResponsibleAI. The 2.6 level of efficacy increase is quite low compared with the potential safety improvements. Obstacle-detection recall rises from 0.88 to 0.97, and the rate of human override of the automatic controls rises from 62.4% to 97.3%.

Similarly, the environment-intervention precision significantly improves from 82.1 to 95.2. The modeled high-risk-event rate significantly reduces from 12.0% under Automation-First operation to 2.6% under Responsible AI, which is an 9.4% absolute reduction and a 78.3% relative reduction.

The Human-in-the-Loop setup proves to be the safety champion, but only by scoring a raw task-efficiency rating of 86.4. Responsible AI makes up for some of this efficiency by allowing for even more autonomous control when faces with a low-risk state but still practicing the aggressive escalation setup.

**Table 3. simulated Logistic Regression predicting high risk Autonomous Field event**


| Predictor                    | B      | SE    | Odds Ratio | 95% CI for OR | p-value |
|------------------------------|--------|-------|------------|---------------|---------|
| Low Visibility               | 1.024  | 0.201 | 2.78       | 1.88–4.11     | <.001   |
| Sensor Disagreement          | 1.230  | 0.223 | 3.42       | 2.21–5.30     | <.001   |
| Worker / Animal Proximity    | 1.322  | 0.237 | 3.75       | 2.36–5.96     | <.001   |
| Network / Communication Loss | 0.652  | 0.191 | 1.92       | 1.32–2.79     | .001    |
| Redundant Sensing            | −0.821 | 0.176 | 0.44       | 0.31–0.62     | <.001   |
| Dynamic Speed Reduction      | −0.673 | 0.169 | 0.51       | 0.37–0.71     | <.001   |
| Human Override               | −1.022 | 0.184 | 0.36       | 0.25–0.52     | <.001   |
| Fail-Safe Stop               | −1.273 | 0.198 | 0.28       | 0.19–0.41     | <.001   |

Table 3 shows the worker/animal proximity id that has the highest positive odds ratio (3.75) and is the greatest modeled predictor of an event with high risk. the sensor disagreement, has a very high odds ratio (3.42), indicating that the conflicting information on perception data is a major safety signal to be treated as an engineered problem.

Low visibility is  $2.78\times$  the modeled odds of a high-risk event, while lost communication has a smaller, yet still significant, impact. These figures demonstrate why an autonomous system must possess degraded-operation notions: keeping full-speed autonomy during sensor disagreement or reduced visibility would subject known uncertainty to operational risk.

The governance variables follow the inverse pattern. Failsafe stopping provided the highest tested modeled protective influence, with a modeled OR of 0.28. Human override lowered the modeled OR to 0.36, redundant sensing to 0.44 and dynamic speed reduction to 0.51. Again, we see evidence that one should not rely on any single safety feature but rather a layered safety architecture.

**Graph 1. Simulated Safety Risk Across Autonomous Agricultural Operating Conditions**



Graph 1 depicts the simulated high- risk-incident rate for each of the 4 field scenarios under the two AI architectures. Operating in open, dry conditions yields the most low risk levels under both architectures. The Automation- First scenario has a high risk incident rate of 4.2%, which drops to just 1.5% under Responsible-AI.

This difference becomes amplified for more complex environments. Similar dense crop rows cause Automation-First rates of 8.6% and 3.1% respectively. Dust or low visibility environments cause the Automation-First rate to reach 15.8% as opposed to Responsible-AI controls which cause a 5.8% rate. The maximum modeled risk occurs when animals or

workers are situated in proximity to the machines operations. Under Automation-First control the rate hits 19.4% while Responsible AI control results in a rate of 6.7%.

The graph shows that governance delivers its maximum absolute gains exactly at the point where complete freedom from interference is least justifiable. Responsible AI doesn't just improve normal operation slightly. It alters how the system operates in the face of challenging environmental circumstances, decelerating machinery, demanding deeper perception consensus, increasing uncertainty, and finally a safe state where appropriate.

## Discussion / Critical Discussion

### a) Representation of the safety and autonomy result.

The results indicate that using the percentage of time without human operation as the measure of autonomy for agricultural machinery can be misleading. This indicator can inadvertently incentivize the analyst to remove safety constraints, and a machine that prevents itself from stopping can seem more autonomous but less believable.

Architecture for Responsible-AI has a different principle: autonomy is defined as a conditional privilege given that evidence justifies safe functioning. In a happy world the machine would keep a large amount of independent power. But as visibility dims, sensors conflict, workers enter the operational zone, communications break down, the safe envelope of autonom....

It accounts for the observed efficiency tendency. Human-in-the-Loop design will prove unoptimal because more human-citations mean more operational interruptions. Responsible AI mitigates the issue because it marks the separation between routine and high-risk events and then human attention is allocated at the place where it is needed rather than a general application.

Then, the decrease from 12.0% to 2.6% modeled occurrence rate of high-risk-events was derived from restriction, not removal, of autonomy, making it an approach more similar to risk-sensitive supervision than to either fully manual or fully independent machine operation.

### b) Comparison With Previous Research

The simulation seems to reflect Bechar and Vigneault's systems perspective that highlights the interdependence of the various sub-systems for sensing, control, mobility and task

performance of agricultural robots. The redundancy of perception in the current model also appears to comply with the work on using the combination of vision, LiDAR and ultrasonic information to achieve agricultural navigation.

The value given to environmental conditions also parallels the more general autonomous-vehicle literature in agriculture. As observed in the review by Roshanianfard and Noguchi, unmanned agriculture vehicles have to function in field conditions where navigation, sensing, and control are still technically difficult. Computer-vision literature also suggests the great promise and reliance of agricultural automation on robust visual sensing.

Furthermore, the 'governance' aspect expands those specific works. In the data, power and distribution, benefits concerned responsible innovation questions do Van der Burg et al. and colleagues; In the digital agriculture, reshaping on those knowledge, skills and institutional relationships in agricultural sector have been observed by Klerkx et al. and colleagues. The present framework has considered the above aspect as autonomous machine is a robot as well as a decision system based on data.

Funding transparency is also consistent with the result that trust in smart farming is linked with transparency and benefit sharing, found by Jakku, et al. An autonomous machine that contributes a clear explanation " *camera confidence drop below threshold*, we decided to stop operation and LiDAR+visual disagree" can afford to be more accountable than a machine that merely offers a generic system fault.

### c) Why Responsible AI Produces a Different Performance Profile

The Responsible-AI architecture works in this way since it does not prioritize maximum task continuity as the primary goal. At each autonomous decision point, there is the chance that continuation would be of lesser value than entering the safe state.

Fail-safe stopping achieves the highest protective coefficient in the simulation. This is especially important as one need not have the machine successfully solve all perception problems in order to proceed with a controlled stop.<sup>12</sup> When the evidence is marginal, the machine can choose not to make a confident physical action based on a precarious perception.<sup>12</sup>

Redundant sensing offers a related benefit. The point of deploying multiple sensors isn't just to get a more accurate average estimate of the world. Their divergence is information itself. If the vision system sees an empty path but the LiDAR finds a nearby obstacle, a sensible

machine should realize its system simply hasn't got sufficient evidence to safely operate normally.

Another way dynamic speed control cuts down the risk is by another method. Agricultural equipment cannot avoid uncertainty at any instant of time, but it can avoid it by reducing speed. As we all know, the reducing of speed, increases response time and generates a longer stopping distance, together with more time for the sensor to double check and for the human to intervene.

Human override serves a supplementing function. Human operators can have background knowledge that isn't encoded in the machine perhaps they know that workers are entering the field, the animals are escaping into an operating area, a sensor is contaminated, or the weather will soon be changing. Oversight thus does not mean that humans outperform AI at continuous perception; it provides contextual judgment and responsibility in making sense of then-machine assumptions.

## Research Contribution / Implications

This analysis yields the most significant theoretical contribution to the concept of ethical autonomy envelopes. Autonomous farm equipment should not be dichotomized into autonomous versus non-autonomous. Rather, permissible machine authority should vary with the combination of environmental risk, perceptual accuracy, task consequence, and human oversight.

For agricultural-machinery manufacturers, this means safety considerations during design evaluation must cover not only successful task achievement, but the system's behavior when it cannot be certain. Safety documentation must specify the operational hazards of sensors, degraded condition modes, stopping performance, available override modes, residual hazards and unacceptable situations for autonomous operation. An important safety precedent for this kind of design consideration can be found in the old ISO 18497:2018.

For farmers and agri-business management, purchasing should answer to the following question rather than 'what is the field capacity'...

- What occurs if the cameras get covered in dust?
- Is this machine capable of distinguishing sensor conflicts?
- Requesting dead putts from the operator. How quickly an operator can regain control.

- Is the system failsafe after Loss of Communication.
- Are critical software updates properly documented?
- What field data are sent out?
- Is autonomous work achievable even when not connected through the cloud?
- Is the decision of the machine and emergency occurrences clearly auditable?
- ‘How do the workers know that the autonomous machine (or to use the context5 ‘the robot’) is running?
- Does the system tell apart persons from other typical obstacles with sufficient safety margin?

Regulators and standards bodies: To enable responsible deployment of agricultural AI, regulators and standards organizations need to ensure that the safe operation of farm machinery, software safety, information security, data management, and human factors are aligned. As shown by the framework for small autonomous agricultural robots developed by Basu et al., it is possible and important for regulatory systems to influence deployment pathways,

If the human role does not become redundant, then automation may well lead to a realignment of roles of an even more profound nature, perhaps resulting in a reduction in the requirement for routine drudgery, such as manual driving and routine treatment application, and a rising need for roles such as monitoring and calibration, machine-maintenance, data analysis, problem solving and safe fleet management. Existing research into digital-agriculture already recognizes that many changes are occurring in fields of knowledge, skills and institutional mechanisms which are adaptations to the technological transformation.

Thus, in the case of environmental governance, the potential for robotic precision is enormous, but only if off-target errors are factored into the measure of performance. An autonomous sprayer should be deemed a success if it can get from A to B fast; precision requires a consideration of whether treatment was applied to the correct plant, at the right location, in the right amount, at the right time.

Methodologically, the paper also demonstrates the possibility of mapping different ethical controls in quantitative terms without claiming that ethics can be translated into a single number. Safety incidents, accuracy of environment, human override, time efficiency and uncertainty are all kept separate.

## Conclusion

The use of autonomous agricultural equipment has the potential to revolutionise field operations, integrating machine perception, autonomous navigation, precision treatment and data-based decision support. The literature on ag-robotics shows considerable advances in autonomous vehicles, robotic field operations, machine vision and multisensor navigation. However, technological capability alone does not guarantee that opened machine autonomy will be ethically feasible.

The hybrid simulation created in this study illustrates that higher raw efficiency does not result in superior total performance. The automation-First tools are slightly better than human-centric tools in field-task efficiency score but are significantly worse in obstacle-detection recall, human-override success, environmental precision, and safety scores. Responsible AI preserves much of the operational efficiency while lowering the modeled high-risk-event rate from 12.0% to 2.6%.

Results further show the context dependence of risk. Dust, reduced visibility, disagreement among sensors, and the presence of humans or animals each significantly led toward increased computed risk. A responsible machine would change its own behavior in response to each one. Decreased speed, redundancy in sensing, human intervention, and fail-safe stopping do not violate the ideal of autonomy; they exemplify controls on autonomy.

An ethical autonomous farm must, then, be a farm that has an adaptive autonomy (not max autonomy). Farmers and engineers should not be striving to achieve the maximal autonomy where it can be achieved technically. What is needed is to position the authority of operations at the best balance between control by humans and control by machines, given the current amount of risk and uncertainty.

Lastly, the responsible use of agricultural robotics should take broader issues into account than just physical safety. Aspects of farm-data governance, cybersecurity, conserving the environment, skills of workers, accountability and long-term reliance on proprietary platforms are also valid elements of technology responsibility. The field of ethics of smart farming and some areas of research within digital agriculture have already proved that processes that inspire transparency, find appropriate levels of trust, control data and ensure the sharing of technologies will better serve responsible innovation. The evaluation of

autonomous machinery should therefore be considered as a sociotechnical agricultural system rather than only as a technological product.

## Limitations

The most serious limitation is all data is synthetic. No real autonomous tractor, field robot, agricultural worker, farm, machinery manufacturer, spraying system, commercial perception model, or accident record was tested.

It must therefore be emphasized that the high-risk-event rates modeled here cannot be quoted as observed probabilities of agricultural accidents. They are intended simply to illustrate the contrasts between alternative governing regimes, using an idealized setting.

The simulation for instance, may be too simple: real farms are more complex than this, as real fields contain a great number and blend of combinations of terrain, crop type, weather conditions, visibility, equipment, animals, people, obstacles and communications properties than the four identified analytical categories.

The logistic-regression coefficients are only simulation results, and are not directly inferred from actual farms.

Human override is also simplified. Actual override success depends on operator training, physical proximity, communication delay, interface design, workload, fatigue, situational awareness, control architecture.

Cyber security was known in concept but not tested through a technical penetration or threat-modeling experiment. Future empirical deployments should separately test for (among other issues) secure communications, software integrity, authentication, remote access, updating, and post-compromise recovery.

The effects of the technology must also be considered beyond this study, such as cost of acquisition, maintenance cost, insurance cost, effect on farm consolidation and technology lock in, and disparities in cost of technology adoption between big and small farms.

## Future Work

Experiments with real robot platforms in test fields should be carried out in future to assess the performance of the Responsible Agriculture Autonomy Framework before implementation in the operational farm environment. Safety experiments should test obstacle-

detection failure in dust, fog, sunlight variation, crop occlusion, dust, mud, sensor failure, communication failure and combinations of human and machine operation.

A second research priority is systematic assessment of human override. For that, researchers should assess detection-to-warning time, operator response time, end-of-stop latency, control transfer success, interface understanding, and post-intervention recovery.

Environmental assessment should be greatly increased. Autonomous spraying, weeding, harvesting and cultivation systems should be evaluated for their impacts to non-target areas, damage done to crops, effect on the soil, effect on non-target organisms, effect of drift, and energy requirements.

Cyber-security research is needed to identify attack traces for the networked cyber-physical systems in autonomous farm machinery. The threat models would include compromised sensors, manipulations of maps, GNSS jamming, unauthorized remote accesses, malicious software upgrades, injection of false obstacles, and denial of communication.

Future work should explore adaptive autonomy limits as well. Instead of using a one-size-fits-all value of confidence, an expert system could learn situational limits that would approach more conservative values as the robot neared a worker, a herd of animals, a road or structure, or any environmentally sensitive area.

To account to the operator, explanations must be empirically tested. Explanations must be judged on their effectiveness to aid safe decision making, not simply their apparent understandability.

Longitudinal studies need to look at the changing role of autonomous machines in agriculture. Some issues to investigate include the role of technicians skills, farmer autonomy, repair rights, use of software, interoperability and access of small farms to technological goods.

Finally, future ethics research should consider responsibility in autonomous incident cases. While technical records may document the machine's beliefs and desires, social institutions are still needed to assign responsibilities among operators, owners, manufacturers, programmers, service providers and data-platform firms.

## Declarations

- **Conflict of Interest:** The fictitious sample author of this sim-based manuscript has stated they have no potential conflict of interest with this work. This declaration is to be replaced by the real author before publishing.
- **Funding:** No external funding is ascribed to this simulation. Any institutional funding, manufacturer support, agricoles-industry sponsorship, supplies of equipment, or research grant related to any future weather-agronomy work should be transparently identified.
- **Ethics Statement:** No real field-operation episodes were involved in this study, only synthetic episodes. There was no human subject or agricultural-worker subject, nor animal subject, nor subject of a farm record containing personally identifiable information, nor subject of an accident record, nor proprietary machine dataset, nor chemical subject of experimental exposure. Thus no human-subject or animal-research approval is necessary for this simulation. All appropriate institutional, safety, and other approvals should be sought for any empirical testing with human subjects, animals, or working-machines.
- **Data Availability:** All numerical results published in this paper are simulation-based methodology data. No confidential commercial or farm dataset has been used.

## References

1. Basu, S., Omotubora, A., Beeson, M., & Fox, C. (2020). Legal framework for small autonomous agricultural robots. In *AI & Society*.
2. Bechar, A. and Vigneault, C. (2016). Agricultural robots for field operations: Concepts and components. *Biosystems Engineering*, 149, 94–111.
3. Bechar, A., & Vigneault, C. (2017). Agricultural robots for field operations. Part 2: Operations and systems. *Biosystems Engineering*, 153, 110–128.
4. Benos, L., Bechar, A., & Bochtis, D. (2020). Safety and ergonomics in human–robot interactive agricultural operations. *Biosystems Engineering*.
5. Chaterji, S., DeLay, N., Evans, J., Mosier, N., Engel, B., Buckmaster, D., & Chandra, R. (2020). Artificial intelligence for digital agriculture at large scale: Strategies, policies, and hurdles.

6. Duckett, T., Pearson, S., Blackmore, S., Grieve, B., Chen, W.-H., Cielniak, G., Cleaversmith, J., Dai, J., Davis, S., Fox, C., et al. (2018). *Agricultural Robotics: The Future of Robotic Agriculture*. UK-RAS Network.
7. European Commission High-Level Expert Group on Artificial Intelligence. (2019). *Ethics Guidelines for Trustworthy AI*. European Commission.
8. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., et al. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28, 689–707.
9. ISO (2018). International organization for standardization:. *18497:2018 Agricultural machinery and tractors Safety of highly automated agricultural machines Principles for design*.
10. Jakku, E., Taylor, B., Fleming, A., Mason, C., Fielke, S., Sounness, C., & Thorburn, P. (2019). If they don't tell us what they do with it, why would we trust them? Trust, transparency and benefit-sharing in smart farming. *NJAS: Wageningen Journal of Life Sciences*.
11. 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.
12. Lioutas, E. D., Charatsari, C., La Rocca, G., & De Rosa, M. (2019). Key questions on the use of big data in farming: An activity theory approach. *NJAS: Wageningen Journal of Life Sciences*.
13. Regan, Á. (2019). “Smart farming” in Ireland: A risk perception study with key governance actors. *NJAS: Wageningen Journal of Life Sciences*.
14. Relf-Eckstein, J. E., Ballantyne, A. T., and Phillips, P. W. B. (2019). Farming reimagined: A case study of autonomous farm equipment and creating an innovation opportunity space for broadacre smart farming. *NJAS: Wageningen Journal of Life Sciences*.
15. Roshanianfard, A., & Noguchi, N. (2020). A review of autonomous agricultural vehicles. *Journal of Terramechanics*.
16. Rovira-Más, F., et al. (2020). Augmented perception for agricultural robots navigation. *IEEE Transactions on Automation Science and Engineering*.

17. Shamshiri, R. R., Weltzien, C., Hameed, I. A., Yule, I. J., Grift, T. E., Balasundram, S. K., Pitonakova, L., Ahmad, D., & Chowdhary, G. (2018). Research and development in agricultural robotics: A perspective of digital farming. *International Journal of Agricultural and Biological Engineering*, 11(4), 1–14.
18. Tian, H., Wang, T., Liu, Y., Qiao, X., & Li, Y. (2020). Application of computer vision technology in the field of agricultural automation a review. *Information Processing in Agriculture*.
19. Van der Burg, S., Bogaardt, M.-J., & Wolfert, S. (2019). Ethics of smart farming, current issues and new directions for responsive innovation to face future challenges New Zealand *Journal of Agricultural ...*
20. Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M.-J. (2017). Big Data in smart farming a review. *Agricultural Systems*, 153, 69 80.