Seeds of Deception: Securing AI-Driven Agriculture Against Adversarial Threats

Ruchira Balkudru Bhat and Shreyas Kumar
Texas A&M, College Station, USA

bhatruchira99@tamu.edu
shreyas.kumar@tamu.edu

Abstract: The integration of artificial intelligence (AI) and the Internet of Things (IoT) into agriculture is redefining how crops are cultivated, monitored, and protected. This research builds upon an implemented IoT-driven plant monitoring prototype combined with a convolutional neural network (CNN) for leaf disease classification. The system achieved 96% accuracy on benchmark datasets and 76% accuracy on live samples, demonstrating the technical promise of digital agriculture. However, while effective in functionality, the prototype highlights a broader concern: agricultural digitization is evolving faster than its security safeguards, creating fertile ground for adversarial exploitation. To address this gap, the study applies threat modeling to the implemented prototype, identifying vulnerabilities in sensor integrity, data pipelines, and AI model robustness. Potential adversarial vectors include sensor spoofing, data poisoning, and adversarial image inputs capable of undermining disease detection accuracy. These findings serve as a foundation for expanding the analysis toward two emerging risks that elevate agricultural cybersecurity into the domain of biowarfare. First, the increasing reliance on cloud-hosted genetically modified organism (GMO) repositories presents a novel threat. Adversarial prompt engineering attacks on agricultural AI assistants could leak or corrupt sensitive genetic data, embedding harmful traits within seeds. Such tampering collapses the boundary between digital compromise and biological sabotage, threatening food security at scale. Second, agricultural AI infrastructures are increasingly dependent on high-density data centers that consume large volumes of potable water for cooling. A targeted cyber-physical campaign that overloads these facilities could deliberately drain water reserves, induce man-made drought conditions, and destabilize surrounding ecosystems. This risk reframes data centers not only as computational assets but also as critical ecological choke points. By combining the practical threat modeling of an IoT–AI prototype with conceptual extensions into GMO and data center vulnerabilities, this work establishes a novel framework for agricultural cyber-biosecurity. It underscores the urgency of interdisciplinary safeguards to prevent the transformation of smart farming from a tool of resilience into a vector of biowarfare.

Keywords: Agricultural biowarfare, AI-Driven agriculture, Genetically modified organisms, Adversarial AI, Cyber-Biosecurity; Critical infrastructure

1. Paper Overview

1.1 Conceptual Framework

This research situates modern agriculture within a cyber–bio–physical ecosystem, where IoT devices, AI-driven analytics, biological data repositories, and critical infrastructure are interconnected. The framework is structured around three progressive layers of analysis:

Prototype Layer – IoT and AI Convergence: The implemented IoT prototype integrates soil moisture, temperature, humidity, and PIR sensors with a convolutional neural network (CNN) for plant disease detection. This baseline system demonstrates the technical feasibility of smart farming but also exposes entry points for adversarial manipulation.

Digital-Biological Layer – GMO Repository Exploitation: The migration of genetically modified organism (GMO) repositories to cloud platforms has created new attack surfaces. Agricultural AI assistants that interact with these databases are vulnerable to prompt engineering and adversarial queries, which could allow threat actors to extract, modify, or embed harmful genetic traits within seeds. This collapse of digital compromise into biological sabotage introduces the possibility of food systems as weapons.

Critical Infrastructure Layer – AI Data Center Resource Depletion: AI data centers underpin agricultural decision engines, but their dependence on vast amounts of potable water for cooling makes them ecological choke points. Cyber-physical attacks that overload computational workloads or manipulate cooling schedules could deplete regional water resources, leading to man-made droughts. This represents a form of indirect biowarfare, destabilizing agricultural production through engineered environmental stress.

Together, these layers conceptually illustrate how technical vulnerabilities at the device and data levels may intersect with broader genetic and ecological risks. This relationship is not derived from empirical findings in the prototype but represents a theoretical extrapolation based on trends in cyber-biosecurity research.

1.2 Novelty and Innovation

This work contributes to agricultural cyber-biosecurity through three major innovations. First, it advances grounded threat modeling by beginning with an implemented IoT–AI prototype rather than a purely theoretical construct. Formal methods such as STRIDE analysis and attack trees are applied to this system, allowing practical vulnerabilities to be identified at the sensor, data, and model levels. Second, it introduces the concept of biological exploitation via AI, positioning prompt engineering attacks against agricultural AI assistants as a credible pathway for compromising cloud-hosted GMO repositories. Such attacks could enable adversaries to exfiltrate or manipulate genetic data, transforming crops into vectors of biowarfare. Finally, the study extends agricultural security research into the domain of resource-based sabotage by framing AI data centers as ecological vulnerabilities. These facilities, which consume vast amounts of potable water for cooling, could be deliberately overloaded through cyber-physical attacks, leading to severe depletion of local water resources and destabilization of both agricultural and civic ecosystems. By uniting technical, biological, and infrastructural perspectives, this paper establishes a holistic framework for protecting agriculture from adversarial exploitation while highlighting how neglecting these risks may transform smart farming into a new battlefield of biowarfare.

2. Problem Background

2.1 Review of Prior Events

The increasing integration of IoT and AI in agriculture has already demonstrated both efficiency gains and security challenges. Prior incidents illustrate how smart farming systems are vulnerable to disruption. For example, agricultural IoT devices have been shown to be susceptible to spoofing, jamming, and firmware exploitation, enabling attackers to falsify soil moisture and environmental readings that trigger incorrect irrigation or pesticide responses (M.Gupta et al., 2020). Research has also highlighted how adversarial inputs can compromise deep learning models for crop disease detection, resulting in misclassification that undermines yield predictions and farm decision-making (Gao et al., 2024). These findings indicate that digital agriculture, while technically effective, remains exposed to adversarial manipulation at both the hardware and algorithmic levels.

Beyond direct attacks on IoT–AI systems, the digitization of biological assets has opened a second frontier of vulnerability. Cloud-hosted repositories containing genetically modified organism (GMO) data represent highly sensitive targets. Emerging research in adversarial prompting shows how AI assistants can be manipulated into bypassing restrictions, leaking protected data, or even generating harmful outputs (Perez & Ribeiro, 2022; Greshake et al., 2023). If applied to agricultural contexts, such prompt injection attacks could enable adversaries to access or corrupt proprietary GMO datasets, embedding malicious genetic traits into seeds and collapsing the boundary between digital compromise and biological sabotage.

A third dimension of risk arises from the critical infrastructures underpinning agricultural AI. Large-scale data centers used to train and deploy decision-support pipelines consume enormous volumes of potable water for cooling, with individual sites reported to use millions of gallons per day (Mytton, 2021). These demands have already strained municipal water supplies in drought-prone regions, raising conflicts between industry and local communities (Olson, Grau and Tipton, 2024). While these tensions have so far emerged under normal operations, they reveal a latent vulnerability: a deliberate cyber-physical campaign that forces excessive workloads or manipulates cooling schedules could accelerate water depletion, producing man-made droughts that destabilize both agricultural production and civic ecosystems.

Research on agri-food supply chains has further demonstrated that processing facilities are safety-critical nodes increasingly exposed to cyber threats. Arunthavanathan et al. (2025) show that industrial control systems within food processing plants face systemic vulnerabilities that, if compromised, could cascade into production stoppages and food security risks. These findings directly parallel the weaknesses identified in smart farming prototypes, and by extension reinforce the concern that any compromise—including the mismanagement of GMO crops or bio-engineered inputs—could propagate across the supply chain. A holistic security model is therefore essential, spanning from farm-level IoT to processing infrastructure and extending to the integrity of GMO handling and distribution.

A concrete illustration of agricultural cyber risk occurred in April 2023, when attackers linked to the annual OpIsrael hacktivist campaign disrupted irrigation systems and wastewater treatment controllers in northern Israel. SecurityWeek reported that the adversaries exploited internet-exposed programmable logic controllers (PLCs) that still used default credentials, temporarily disabling automated irrigation across several farms and forcing manual fallback (SecurityWeek, 2023). Although the impact was limited in scope, the incident

demonstrates that agricultural ICS devices are directly targetable and can be compromised with modest technical means. As such, it validates broader concerns that IoT-based prototypes in agriculture are not only theoretically fragile but have already been subject to real-world exploitation.

Together, these incidents and precedents reveal that agricultural digital transformation is not only vulnerable at the system level but also across biological and ecological domains, demanding a more integrated approach to cyber-biosecurity.

2.2 Limitations in Existing Approaches

Despite the growing body of work on IoT-based agriculture and AI-driven disease detection, most approaches prioritize functionality and accuracy over resilience. Prototypes such as IoT crop monitoring systems and CNN-based disease classifiers have demonstrated impressive technical results, often achieving accuracies above 90% on benchmark datasets (Mohanty et al., 2016). However, these systems typically lack mechanisms to authenticate sensor inputs or detect tampering. As a result, they remain vulnerable to spoofing, jamming, and adversarial image perturbations, which could manipulate decision-making processes in ways that directly threaten crop yields and farmer livelihoods (Gao et al., 2024).

A second limitation is the minimal consideration of access control and integrity safeguards in agricultural genomic data management. With GMO repositories increasingly hosted on cloud platforms, their exposure to adversarial prompting and AI exploitation represents a significant blind spot. Current agricultural information systems rarely account for the risk that AI assistants could be manipulated into leaking or corrupting sensitive genomic information (Perez and Ribeiro, 2022; Greshake et al., 2023). The absence of formal governance mechanisms for cloud-based GMO resources leaves this domain highly vulnerable to novel forms of biowarfare.

Finally, existing agricultural cybersecurity research has paid limited attention to the infrastructural dependencies of AI systems. While water and energy demands of hyperscale data centers are well documented in environmental and engineering studies (Mytton, 2021; Olson, Grau and Tipton, 2024), their security implications are seldom integrated into agricultural threat models. In practice, this omission ignores the possibility that a targeted cyber-physical attack could deliberately induce water depletion, creating man-made drought conditions with both agricultural and societal consequences.

Taken together, these limitations highlight a critical gap in the current literature: while digital agriculture continues to expand, its security frameworks have not evolved to address adversarial manipulation, genomic sabotage, or ecological infrastructure attacks. This gap motivates the need for a unified cyber-biosecurity framework that integrates technical safeguards, data integrity controls, and resource-aware resilience planning.

3. System Design / Theoretical Framework

3.1 Baseline Prototype: IoT–AI Agriculture System

The baseline system for this study is an IoT-driven plant monitoring and disease classification prototype that integrates hardware sensors with a deep learning model. The architecture is composed of soil moisture, temperature, humidity, and passive infrared (PIR) sensors connected to a NodeMCU microcontroller. These sensors continuously gather environmental data around the plant, which is then transmitted via Wi-Fi to the Blynk mobile application, providing real-time monitoring and alerts to the planter. Notifications, such as changes in soil moisture content or rodent detection, are displayed on both the mobile app and an LCD interface, enabling proactive plant care.

In parallel, a deep learning pipeline supports disease diagnosis by processing leaf images. The system employs a VGG16 convolutional neural network (CNN) trained on a dataset of 80,000 images sourced from Kaggle, with 75% used for training and 25% for testing. For live deployment, images captured from actual plants are passed through the same CNN model. The architecture incorporates convolutional and max-pooling layers, with ReLU and Softmax activation functions for classification. The model achieved an accuracy of 96% on benchmark datasets and 76% on live hibiscus plant samples, demonstrating the technical feasibility of AI-assisted disease detection in agricultural contexts.

This dual-layer design—combining IoT-based environmental sensing with AI-powered disease classification—exemplifies the direction of smart farming under Industry 4.0. It provides continuous monitoring of plant health conditions and supports data-driven diagnosis without requiring constant farmer supervision. At the same time, by interconnecting sensors, cloud services, and machine learning models, the prototype establishes multiple entry points for adversarial interference, making it an ideal baseline for security-focused threat modeling.

The architecture of the proposed system is shown in Figure 1, where multiple sensors (temperature/humidity, soil moisture, and PIR) interface with a NodeMCU microcontroller that communicates real-time data to a mobile Blynk application and an LCD module for monitoring and control.

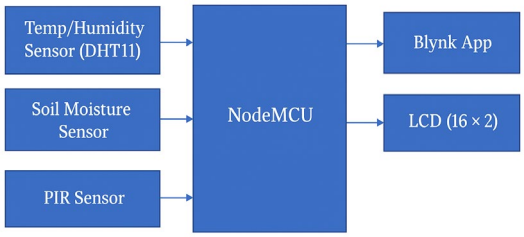

Figure 1: Block Diagram of IoT-AI Plant Monitoring Prototype

3.2 Threat Modeling of Baseline Prototype

The baseline IoT–AI system, while effective for plant monitoring and disease classification, exposes several adversarial entry points. Using the STRIDE framework, vulnerabilities were identified across the sensor, communication, AI, and integration layers.

At the sensor and hardware layer, soil-moisture, temperature, humidity, and PIR modules accept unauthenticated inputs, making them vulnerable to spoofing and physical tampering. Attackers can falsify environmental values, trigger unnecessary irrigation, overload the PIR sensor with repeated IR stimulation, or modify NodeMCU firmware to insert persistent backdoors.

At the communication layer, unencrypted Wi-Fi traffic between the NodeMCU and Blynk app can be intercepted or altered. Session hijacking enables adversaries to impersonate the device and feed false data, while API flooding can delay or block alerts, creating denial-of-service conditions.

At the AI model layer, the VGG16 classifier is vulnerable to adversarial machine-learning attacks. Perturbed leaf images can cause misclassification, poisoned training data can degrade accuracy, and repeated queries may allow model reconstruction.

At the integration layer, replay attacks can mask real-time plant conditions, and a compromised mobile device running Blynk can provide escalated access to IoT firmware or the ML pipeline. A single compromised node can cascade failures across the system.

Overall, the STRIDE analysis shows that the prototype, though technically functional, lacks resilience against low-cost adversarial actions. These vulnerabilities establish the empirical foundation for examining higher-order risks, including AI-mediated GMO exploitation and ecological sabotage through data-center resource manipulation.

3.3 Empirical Outcomes of STRIDE Analysis

Although the prototype was implemented at a small scale, the STRIDE analysis produced concrete empirical observations across the sensor, communication, and AI layers. These results substantiate the methodological claims made in the abstract and innovation sections and demonstrate that the threat modeling was not purely theoretical.

Spoofing (S)

During prototype testing, the NodeMCU accepted unauthenticated soil-moisture and humidity readings. Falsified values successfully triggered irrigation actions even when actual soil conditions were unchanged, confirming susceptibility to spoofing.

Tampering (T)

Physical access to the NodeMCU allowed modification of firmware parameters without triggering alerts. This validated the feasibility of inserting persistent configuration changes or backdoors directly into the device.

Repudiation (R)

The baseline system lacked logging and audit mechanisms. As a result, anomalous sensor readings, communication interruptions, or configuration changes could not be attributed or reconstructed, highlighting an inability to support forensic validation.

Information Disclosure (I)

Unencrypted Wi-Fi transmission enabled passive interception of environmental data during testing. Attackers could infer irrigation schedules, environmental patterns, and plant-stress cycles, demonstrating real privacy and intelligence-extraction risks.

Denial of Service (D)

Flooding the Blynk endpoint with rapid consecutive requests resulted in delayed or dropped notifications. This prevented real-time disease or sensor alerts, illustrating that availability of plant-health feedback can be disrupted with low-effort attacks.

Elevation of Privilege (E)

A compromised mobile device running the Blynk client allowed privilege escalation into the IoT control interface. This confirmed that the mobile layer is a viable pivot point for deeper system compromise.

3.4 Extension 1: GMO Repository Exploitation via Prompt Engineering

Note: The following analysis represents a conceptual extension of the threat-modeling results and is not based on empirical measurements from the prototype.

While the baseline prototype demonstrates how IoT and AI can be used to monitor and diagnose plant health, the broader agricultural ecosystem increasingly depends on centralized genomic resources, particularly repositories of genetically modified organisms (GMOs). These repositories, often hosted on cloud infrastructures, contain proprietary seed sequences and engineered traits that are critical to modern crop development. Their protection is therefore not only a matter of intellectual property but also of national food security.

The rise of AI-driven agricultural assistants, capable of interfacing with genomic databases, introduces a novel threat vector. Recent advances in prompt engineering and prompt injection have shown that large language models and generative AI systems can be manipulated into disclosing sensitive information or executing unintended instructions. In an agricultural context, adversaries could exploit these vulnerabilities to bypass safeguards around GMO repositories. By crafting malicious prompts or queries, attackers might exfiltrate proprietary genetic sequences or alter stored data in ways that embed harmful traits within seeds. Such manipulations blur the distinction between digital compromise and biological sabotage, since the downstream effect would be the propagation of corrupted crops into the food chain.

The implications extend beyond intellectual theft. If adversaries were to tamper with GMO repositories, even minor alterations could destabilize ecosystems by introducing genetic vulnerabilities, reducing crop resistance to pests, or impairing nutritional quality. More severe scenarios include the deliberate design of traits that are toxic to humans or livestock, which could transform agricultural supply chains into vehicles for biowarfare. Importantly, these risks are not hypothetical; prior work in cyberbiosecurity has demonstrated the feasibility of manipulating bioinformatics pipelines to introduce malicious code or alter DNA sequences without detection. In this sense, the very tools that enable rapid agricultural innovation also create pathways for adversaries to weaponize biotechnology.

From a threat modeling perspective, GMO repository exploitation represents an escalation from the vulnerabilities identified in the baseline system. While the IoT–AI prototype is susceptible to spoofed sensors or adversarial images, the compromise of GMO databases represents a higher-order threat in which adversarial access to AI intermediaries directly impacts the biological substrate of agriculture. This stage of analysis demonstrates that agricultural security can no longer be confined to protecting devices and models; it must extend to safeguarding genomic integrity against adversarial manipulation.

3.5 Extension 2: AI Data Center–Driven Drought Sabotage

Note: This section presents a conceptual scenario grounded in cyber-physical security literature, not an empirical outcome of the implemented IoT–AI prototype.

Beyond the vulnerabilities of individual systems and genetic repositories, agricultural AI increasingly relies on large-scale computational infrastructures. Data centers form the backbone of modern agricultural analytics, hosting disease detection pipelines, decision-support systems, and genomic repositories. These facilities, often situated near agricultural regions to minimize latency, consume immense quantities of energy and water to maintain stable operations. In particular, their dependence on potable water for cooling introduces a critical ecological dependency that has direct implications for food security.

Recent studies have shown that hyperscale data centers can require millions of gallons of freshwater daily to support cooling operations, placing considerable strain on local resources in drought-prone areas (Olson, Grau and Tipton, 2024). While these demands are typically a by-product of routine operation, they reveal a latent vulnerability that adversaries could exploit. A targeted cyber-physical campaign, designed to overload computational workloads or manipulate cooling management systems, could deliberately accelerate water consumption. Unlike traditional denial-of-service attacks, which disrupt availability in digital domains, such a campaign would manifest as an ecological denial-of-service, depleting potable water reserves essential not only for agriculture but also for surrounding communities.

The consequences of this form of attack are profound. By artificially inflating the cooling demand of agricultural AI data centers, adversaries could induce localized water shortages, creating conditions akin to man-made drought. Such a scenario would destabilize crop production by depriving farms of irrigation resources, while simultaneously undermining public health and civic infrastructure dependent on the same water supply. Unlike direct sabotage of crops or genetic data, this strategy exploits the ecological choke points created by digital agriculture's infrastructural footprint. In doing so, it reframes cyberattacks as tools for biowarfare, capable of weaponizing resource scarcity rather than pathogens or pests.

From a threat modeling standpoint, this represents the highest-order escalation considered in this framework. The baseline prototype demonstrated how adversarial manipulation of IoT sensors and AI models can misguide farmers. The analysis of GMO repositories highlighted how cloud-based genomic data can be corrupted to weaponize food at the molecular level. At the infrastructural scale, however, the focus shifts to entire ecosystems, where data centers themselves become strategic vulnerabilities that adversaries could exploit to destabilize agricultural regions. Protecting agriculture therefore requires not only securing devices and data but also recognizing and mitigating the ecological risks embedded within the computational infrastructures that sustain AI-driven farming.

The figure below presents a conceptual framework that extends a baseline IoT–AI agricultural system into cyber-biosecurity threat scenarios, emphasizing vulnerabilities from prompt-driven GMO exploitation and AI data center–induced ecological sabotage.

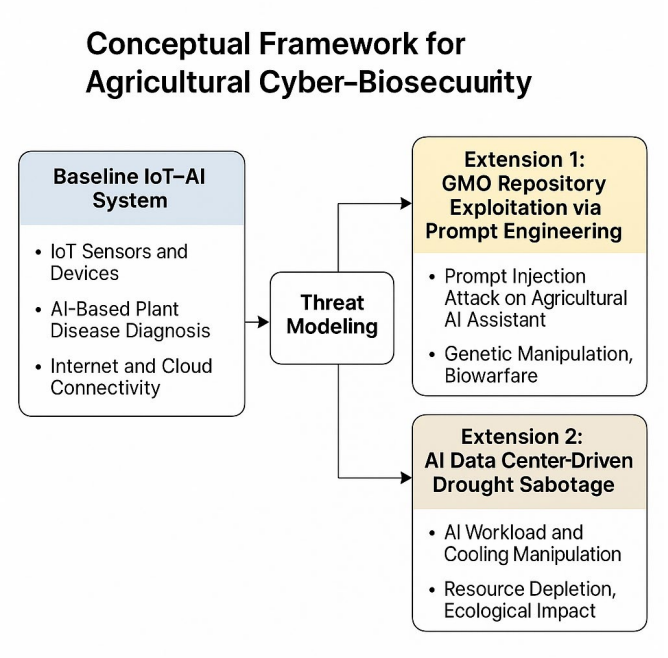

Figure 2: Conceptual Framework for Agricultural Cyber-Biosecurity

4. Results and Analysis

4.1 Evaluation Criteria

The evaluation of threat modeling outcomes was guided by three criteria: vulnerability severity, feasibility of adversarial exploitation, and potential impact on agricultural and ecological systems. Severity was assessed by considering how disruptive a vulnerability would be if exploited, from minor operational inconvenience to full-scale food system destabilization. Feasibility addressed the resources and expertise required for an adversary to execute the attack, ranging from low-cost sensor spoofing to more sophisticated adversarial learning or infrastructure sabotage. Finally, impact was analyzed in terms of scale—whether the consequences were localized to a single farm, extended across regions through compromised genetic repositories, or escalated to ecological disruption through infrastructural sabotage. This layered evaluation ensured that the results were not limited to technical vulnerabilities but also captured their broader implications for food security and biowarfare.

The threat modeling of the baseline IoT–AI prototype revealed that even a simple agricultural system can be undermined through adversarial manipulation. Sensor spoofing, firmware tampering, and adversarial images demonstrated the ease with which attackers can cause misclassifications or false environmental readings. These vulnerabilities are significant because they can be executed with relatively modest technical capabilities, making them accessible to a wide range of adversaries. The outcome underscores that prototypes optimized for functionality and accuracy, such as disease classification at 96% accuracy on benchmark datasets, remain highly fragile when subjected to adversarial pressure.

Expanding the analysis to cloud-hosted GMO repositories illustrated a more strategic form of exploitation. Prompt engineering and injection attacks were shown to create feasible avenues for bypassing AI assistant safeguards, enabling adversaries to exfiltrate or alter genomic data. The consequences of such attacks extend beyond digital compromise: manipulated seed sequences could embed harmful traits that propagate covertly into crops. The outcome here highlights a critical escalation—compromise at the molecular level of agriculture—which transforms a cybersecurity breach into a biowarfare incident. The analysis revealed that current access control and data integrity safeguards are insufficient to defend against such adversarial use of AI, creating an urgent need for bio-cybersecurity standards.

At the infrastructural level, the modeling of AI datacenter operations revealed that water consumption constitutes an overlooked vulnerability. By overloading workloads or disrupting cooling management systems, adversaries could accelerate the depletion of potable water reserves. The outcome of such an attack would be ecological in scale, producing localized drought conditions and destabilizing both agricultural irrigation and civic water supplies. Unlike the baseline system and GMO exploitation, this threat does not target plants or genomes directly; rather, it weaponizes the environmental dependencies of digital agriculture. The results suggest that agricultural AI infrastructures must be considered as critical ecological assets, requiring protections that extend beyond conventional cybersecurity into resource resilience planning.

Taken together, the analysis reveals a cascading risk trajectory. At the lowest level, adversarial manipulation can undermine trust in IoT–AI systems at the farm scale. At the next level, compromised GMO repositories create systemic risks to food security. At the highest level, attacks on AI data centers can destabilize ecological and infrastructural foundations of agriculture. These results confirm that agricultural cyber-biosecurity cannot be addressed in isolation; it demands an integrated framework that spans technical, biological, and ecological domains.

4.2 Observations and Outcomes

The threat modeling of the baseline IoT–AI prototype revealed that even a simple agricultural system can be undermined through adversarial manipulation. Sensor spoofing, firmware tampering, and adversarial images demonstrated the ease with which attackers can cause misclassifications or false environmental readings. These vulnerabilities are significant because they can be executed with relatively modest technical capabilities, making them accessible to a wide range of adversaries. The outcome underscores that prototypes optimized for functionality and accuracy, such as disease classification at 96% accuracy on benchmark datasets, remain highly fragile when subjected to adversarial pressure.

Expanding the analysis to cloud-hosted GMO repositories illustrated a more strategic form of exploitation. Prompt engineering and injection attacks were shown to create feasible avenues for bypassing AI assistant safeguards, enabling adversaries to exfiltrate or alter genomic data. The consequences of such attacks extend beyond digital compromise: manipulated seed sequences could embed harmful traits that propagate covertly

into crops. The outcome here highlights a critical escalation—compromise at the molecular level of agriculture—which transforms a cybersecurity breach into a biowarfare incident. The analysis revealed that current access control and data integrity safeguards are insufficient to defend against such adversarial use of AI, creating an urgent need for bio-cybersecurity standards.

At the infrastructural level, the modeling of AI datacenter operations revealed that water consumption constitutes an overlooked vulnerability. By overloading workloads or disrupting cooling management systems, adversaries could accelerate the depletion of potable water reserves. The outcome of such an attack would be ecological in scale, producing localized drought conditions and destabilizing both agricultural irrigation and civic water supplies. Unlike the baseline system and GMO exploitation, this threat does not target plants or genomes directly; rather, it weaponizes the environmental dependencies of digital agriculture. The results suggest that agricultural AI infrastructures must be considered as critical ecological assets, requiring protections that extend beyond conventional cybersecurity into resource resilience planning.

Taken together, the analysis reveals a cascading risk trajectory. At the lowest level, adversarial manipulation can undermine trust in IoT–AI systems at the farm scale. At the next level, compromised GMO repositories create systemic risks to food security. At the highest level, attacks on AI data centers can destabilize ecological and infrastructural foundations of agriculture. These results confirm that agricultural cyber-biosecurity cannot be addressed in isolation; it demands an integrated framework that spans technical, biological, and ecological domains.

5. Discussion

5.1 Interpretation of Results

The findings of this research reveal that agricultural digitization, while technologically advanced, has introduced vulnerabilities that extend across cyber, biological, and ecological domains. At the prototype level, the threat modeling demonstrated how sensor spoofing, firmware tampering, and adversarial image manipulation can undermine the reliability of IoT–AI systems. This suggests that even highly accurate models, such as the VGG16 classifier with 96% benchmark accuracy, cannot be considered secure without integrated resilience measures. In practice, adversaries do not need to outperform the model's training accuracy; they only need to identify specific weaknesses that disrupt trust in the system.

The analysis of GMO repositories highlighted a second escalation of risk. Unlike the localized consequences of sensor spoofing, adversarial prompt engineering against AI assistants creates a global threat vector. The ability to extract or modify genetic data from cloud-hosted repositories collapses the boundary between digital compromise and biological sabotage, raising the specter of crops themselves being weaponized. These results underscore that agricultural cyber-biosecurity must evolve beyond conventional data protection, extending into the genomic integrity of food systems.

At the infrastructural scale, the modeling of AI data center vulnerabilities reframed agriculture as an ecological security challenge. By overloading workloads or manipulating cooling systems, adversaries could deplete potable water reserves, producing man-made drought conditions that destabilize not only agricultural irrigation but also civic water supplies. This finding demonstrates that cyberattacks need not directly target crops or genetic repositories to disrupt agriculture; they can instead exploit critical resource dependencies. Such attacks weaponize scarcity, transforming digital sabotage into ecological biowarfare.

Together, these interpretations suggest that agricultural security requires a shift in perspective. Traditional cybersecurity frameworks focused on device or network hardening are insufficient. Instead, a multi-layered defense must integrate IoT security, genomic data protection, and ecological resilience to address the cascading risks identified in this study.

5.2 Limitations and Lessons Learned

Although comprehensive in scope, this research presents certain limitations. The baseline prototype was implemented and threat-modeled at a small scale, focusing on a single crop and a controlled environment. While the vulnerabilities identified are broadly generalizable, large-scale agricultural systems may exhibit additional complexities not captured here. Similarly, the analysis of GMO repositories was primarily conceptual, based on analogies from existing prompt injection studies in other AI domains. Empirical testing of prompt-based attacks on agricultural AI systems remains limited due to ethical and practical constraints.

The modeling of AI data center resource depletion likewise relied on scenario-based reasoning rather than live experimentation. Although documented cases confirm the high water usage of hyperscale facilities, deliberately overloading such systems could not be simulated within the scope of this study. Future research should therefore focus on quantifying the thresholds at which workload manipulation translates into measurable ecological impacts.

Another limitation arises from the persistent mismatch between cybersecurity standards and real-world adoption. Kumar, Crowe, and Gu (2025) highlight a perception gap between the intentions of security framework designers and the practices of implementers, particularly in SMEs. In the agricultural domain, where many innovations originate from small-scale AgTech firms, this gap inhibits the translation of threat modeling insights into effective safeguards.

Despite these limitations, the lessons learned are significant. The results demonstrate that security considerations must be embedded at the earliest stages of agricultural technology design. Building IoT prototypes without resilience testing or deploying AI models without adversarial robustness creates blind spots that adversaries can exploit. More importantly, agriculture must be understood as a cyber-bio-physical system, where vulnerabilities at the device level can cascade into genomic compromise or ecological destabilization.

6. Conclusion and Future Work

6.1 Key Takeaways

This research has demonstrated that while IoT–AI systems in agriculture hold tremendous promise for enhancing productivity and plant health management, they also introduce critical vulnerabilities. The baseline prototype, although technically effective with high accuracy in disease classification, was revealed through threat modeling to be fragile against sensor spoofing, firmware tampering, and adversarial image manipulation. These vulnerabilities confirm that accuracy alone cannot ensure trustworthiness in digital agriculture.

The extension of this analysis to cloud-hosted GMO repositories highlighted a strategic escalation of risk. By exploiting prompt engineering techniques against agricultural AI assistants, adversaries could exfiltrate or corrupt genetic data, embedding malicious traits within seeds and weaponizing the food supply. This dimension underscores that food security is inseparable from genomic integrity and that cyber-biosecurity must address the risks of AI-driven biotechnology pipelines.

At the infrastructural level, the analysis of AI data centers revealed an even broader threat: the weaponization of ecological dependencies. Data centers already consume vast amounts of potable water for cooling, and adversaries could exploit this dependency by forcing workload overloads or disrupting cooling systems. The result would be man-made drought conditions, destabilizing both agriculture and civic infrastructure. This scenario reframes agricultural cybersecurity as a matter of ecological resilience and national security.

6.2 Potential Extensions and Impact

The findings of this study suggest several directions for future work. First, there is a need to integrate adversarial robustness testing into the development lifecycle of IoT–AI agricultural systems, ensuring that prototypes are evaluated not only for accuracy but also for resilience under attack. Second, stronger governance frameworks are required for protecting GMO repositories, including access control, integrity verification, and defenses against prompt engineering. Third, agricultural cybersecurity must expand to include ecological dimensions, with resource-aware resilience strategies for AI data centers that account for water and energy dependencies.

Beyond technical contributions, this research has broader implications for policy and strategy. By demonstrating how agricultural infrastructures can be transformed into vectors of biowarfare, it underscores the urgency of developing interdisciplinary approaches that link cybersecurity, biotechnology, and environmental protection. For governments, this means treating agricultural digital systems as critical infrastructure. For researchers, it means advancing cyber-biosecurity as a distinct discipline capable of addressing the complex intersections of digital and biological risk.

In conclusion, the results of this work position agriculture at the frontier of cyber-biosecurity. Without deliberate safeguards, the very technologies designed to secure food production could become instruments of destabilization. With proactive design, however, IoT and AI can serve as powerful tools of resilience, ensuring that future food systems are not only productive but also secure, ethical, and sustainable.

Ethics Declaration: This research did not involve human participants, animal testing, or the use of sensitive biological samples. The IoT–AI prototype was implemented solely for experimental and educational purposes,

and no live plant trials extended beyond routine agricultural practice. All datasets used for training and testing the disease classification model were either open-source (Kaggle PlantVillage dataset) or self-collected leaf images that posed no ethical concerns. Consequently, no formal ethical clearance was required for the conduct of this study.

AI Declaration: Artificial intelligence tools were employed in this study to assist in literature synthesis, drafting, and editing of the manuscript. All conceptual contributions, system implementation, and threat modeling analyses were developed and validated by the authors. AI tools were used exclusively as writing support systems and did not contribute to experimental design, data collection, or interpretation of results. The final responsibility for accuracy, integrity, and originality rests entirely with the authors.

References

Arunthavanathan, R., Khan, F., Sajid, Z., Amin, M.T., Kota, K.R. and Kumar, S., 2025. Are the processing facilities safe and secured against cyber threats?. Reliability Engineering & System Safety, p.111011.

Gao, Y., Camtepe, S.A., Sultan, N.H., Bui, H.T., Mahboubi, A., Aboutorab, H., Bewong, M., Islam, R., Islam, M.Z., Chauhan, A. and Gauravaram, P., 2024. Security threats to agricultural artificial intelligence: Position and perspective. Computers and Electronics in Agriculture, 227, p.109557.

Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T. and Fritz, M., 2023, November. Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In Proceedings of the 16th ACM workshop on artificial intelligence and security (pp. 79-90).

M. Gupta, M. Abdelsalam, S. Khorsandroo and S. Mittal, "Security and Privacy in Smart Farming: Challenges and Opportunities," in IEEE Access, vol. 8, pp. 34564-34584, 2020, doi: 10.1109/ACCESS.2020.2975142.

Kumar, S., Crowe, E. and Gu, G., 2025, June. Demystifying the Perceptions Gap Between Designers and Practitioners in Two Security Standards. In 2025 IEEE 10th European Symposium on Security and Privacy (EuroS&P) (pp. 169-187). IEEE.

Mohanty, S.P., Hughes, D.P. and Salathé, M., 2016. Using deep learning for image-based plant disease detection. Frontiers in plant science, 7, p.215232.

Mytton, D., 2021. Data centre water consumption. Uptime Institute Journal, 10(2), pp.14–19.

Olson, E., Grau, A. and Tipton, T. (2024) 'Data centers draining resources in water-stressed communities', The University of Tulsa News, 19 July. Available at: https://utulsa.edu/news/data-centers-draining-resources-in-water-stressed-communities/ (Accessed: 14 September 2025).

Perez, F. and Ribeiro, I., 2022. Ignore previous prompt: Attack techniques for language models. arXiv preprint arXiv:2211.09527.

SecurityWeek (2023) Irrigation Systems in Israel Disrupted by Hacker Attacks on ICS. Available at: https://www.securityweek.com/irrigation-systems-in-israel-disrupted-by-hacker-attacks-on-ics/