
AI-Safety–Driven System for Predicting Cross-Pollination Risk and Optimizing GMO Testing in Soybean Fields¹

Roberto Nuñez
HTW Saar, Germany

Benjamín López
HTW Saar, Germany

Małgorzata Komorowska
HTW Saar, Germany

Liting Zhang
HTW Saar, Germany

Juan Diego Cordoba
HTW Saar, Germany

Abstract

Unintended cross-pollination between GMO and non-GMO crops poses a significant biosecurity and economic risk, potentially leading to certification loss, rejected shipments, and regulatory interventions. To address this challenge, we present **BioSecure AI**, an AI-safety–driven system that predicts cross-pollination risk and optimizes field testing strategies. The system models adjacent fields using a biologically informed contamination framework combining exponential distance decay, wind-driven dispersal, and structural noise to simulate realistic pollen drift. A hybrid inspection agent—integrating a Deep Q-Network (DQN) with heuristic cluster-skipping—selects sampling locations based on expected return on investment (ROI), balancing testing costs against the risk of undetected contamination.

Results show that BioSecure AI consistently identifies high-risk zones along GMO-adjacent field edges and generates efficient, non-redundant sampling paths. The agent terminates testing automatically when additional actions no longer provide positive expected value, achieving substantial cost savings while maintaining comprehensive risk coverage. This work demonstrates how cybersecurity concepts such as threat mapping, risk scoring, and autonomous mitigation can be applied to agricultural

¹ Research conducted at the Defensive Acceleration Hackathon, 2025

biosecurity, offering a scalable and interpretable tool for safeguarding non-GMO supply chains.

Keywords: *Biosecurity modeling, AI safety, reinforcement learning, contamination risk assessment, cyber-physical monitoring*

1. Introduction

Modern agriculture faces a growing challenge in preventing unwanted cross-pollination between GMO (genetically modified organism) and non-GMO crops, a risk that can compromise certification, trigger contract penalties, and create significant economic losses for farmers. Because pollen drift is influenced by factors such as wind, field geometry, crop species and flowering overlap, contamination events are difficult to predict and even harder to detect early. Our project addresses this vulnerability by developing an AI-driven monitoring system capable of modeling two adjacent fields—one GMO and one non-GMO—and generating a risk heat map that identifies where cross-breeding is most likely to occur. By recommending optimal genetic testing locations and frequencies, the system provides farmers with a low-cost early-warning tool and documented evidence of due diligence. This reduces the probability of contamination going unnoticed and strengthens protection against commercial disputes or legal claims, offering a scalable and accessible way to safeguard non-GMO supply chains, control, and modify their local cyber-security environment in new and changing ways.

However, preventing GMO–non-GMO cross-pollination has been met with a range of practical and economic challenges. Firstly, agricultural environments are highly variable and far less controlled than laboratory or industrial settings. Wind conditions, landscape geometry, crop height, flowering synchrony, and microclimate differences introduce uncertainty that farmers cannot easily model or anticipate. Even highly trained agronomists often lack the real-time, field-level visibility required to identify where contamination risk is concentrated or how it evolves during the growing season.

Secondly, while genetic testing technologies exist, relying on them alone as a detection strategy is both costly and inefficient. Rapid strip assays, ELISA tests, and PCR analyses each carry operational expenses—ranging from USD 20 to USD 100 per sample—and these costs accumulate quickly when farmers attempt to monitor large field perimeters or run tests at high frequency. In contrast, regulatory consequences from a single unnoticed contamination event—such as loss of non-GMO certification, rejected export lots, or intensified inspections—can incur financial losses in the tens or even hundreds of thousands of dollars. Yet these severe penalties, although costly, occur far less frequently than the routine burden of seasonal testing. This discrepancy creates a structural challenge: farmers must

invest disproportionately in prevention without any guarantee that such investment will avert rare—but economically devastating—events.

Therefore, this project focuses chiefly on the democratization of risk-aware monitoring and decision-making in agricultural environments, enabling farmers to access predictive capabilities traditionally limited to research institutions or regulatory bodies. By integrating physics-based pollen-drift modeling with machine learning, the system generates real-time spatial risk heat maps and recommends cost-efficient testing strategies. The reward function explicitly incorporates both low-frequency/high-impact regulatory penalties and high-frequency/low-impact testing costs, allowing the AI to balance proactive risk reduction with long-term operational sustainability. This makes it possible for farmers with no technical background to understand where and when to test, how frequently, and why—effectively providing an accessible early-warning system tailored to real-world agricultural constraints.

The selection of biological parameters in the model directly reflects the empirical behavior of soybean pollen drift and supports the system’s ability to generate realistic and actionable risk predictions. The maximum cross-pollination rate is set between 3% and 8%, representing the upper contamination levels observed when non-GMO soybeans grow immediately adjacent to GMO fields under ideal dispersal conditions. This establishes a biologically defensible ceiling for edge-level contamination. To capture the influence of wind—a dominant but highly variable driver of pollen movement—the model adopts a wind-speed saturation threshold of 5.5 m/s, corresponding to regional averages in Alabama. Speeds above this threshold do not meaningfully increase pollen transport, ensuring the model remains sensitive to weather variation without exaggerating its effects.

To represent how contamination probability diminishes with distance, the exponential decay parameter k is calibrated using the known biological range of soybean pollen travel, typically between 3 and 38 meters. By selecting k such that risk declines to biologically negligible levels beyond 20–50 meters, the model accurately reproduces the sharp spatial attenuation observed in field studies. Together, these parameter choices create a framework that is both scientifically grounded and computationally efficient, enabling the AI system to simulate realistic contamination dynamics while remaining interpretable for farmers and policy stakeholders. This alignment between biological evidence and model structure is essential for producing credible risk maps and reliable testing recommendations within the real-world constraints described above.

2. Methods

Overview

The BioSecure AI project was designed as an autonomous contamination-containment assistant capable of analyzing large agricultural fields, detecting high-risk dispersion zones, and optimizing the deployment of biological test kits using reinforcement learning.

The system was implemented as a complete pipeline consisting of:

- A Flask backend API
- A Streamlit command interface
- A custom risk-field generator
- A Deep Q-Network (DQN) model
- A full inspection strategy combining RL + heuristics

All code modules were structured into small, reusable components and kept in version-controlled GitHub repositories throughout development.

Heatmap Generation Method

The project required a realistic model of contamination spread. Several approaches were experimented:

2.1 Pure exponential decay

The first version used:

$$p(x) = p_0 e^{-kx}$$

This produced a perfect right-to-left gradient.

2.2 Structural noise

To break the unrealistic “solid red column”, controlled noise was added:

```
ruido = np.random.uniform(-0.05, 0.05, (height, width))  
mapa_final = riesgo_base + ruido
```

Image 1 – controlled noise fragment

This created a realistic biological texture.

2.3 Stable synthetic fields for debugging

To help debugging the DQN, deterministic version was added:

```
def generate_heatmap_matrix(height, width):

    p0 = 0.98
    k = 0.005

    Y, X = np.meshgrid(np.arange(height), np.arange(width), indexing='ij')
    distancia_x = (width - 1) - X

    riesgo_base = p0 * np.exp(-k * distancia_x)

    ruido = np.random.uniform(-0.02, 0.02, (height, width))

    return np.clip(riesgo_base + ruido, 0.0, 1.0)
```

Image 2 – heat map matrix generated

This matrix became the core input for all agents.

3. Backend Architecture (Flask API)

The backend exposes a single endpoint:

```
@app.route('/analyze', methods=['POST'])
def analyze_field():
```

Image 3 – post method fragment

It:

1. Receives either:
 - dimensions (width, height), or
 - a custom risk map
2. Generates or loads the contamination matrix
3. Feeds it to the DQN agent
4. Runs a multi-step mission planner
5. Returns JSON containing:
 - selected coordinates
 - confidence per scan
 - ROI

- stop condition
- full mission plan

Key α -algorithm

The backend uses hybrid greedy + RL:

```
max_pos = np.unravel_index(np.argmax(current_map), current_map.shape)
center_y, center_x = max_pos
```

Image 4 – search for coordinates fragment

Instead of scanning all cells in a line, the system jumps between clusters by clearing a 150-pixel AoE:

```
current_map[y_start:y_end, x_start:x_end] = 0.0
```

Image 5 – cleaning neighbourhood fragment

This prevents the agent from repeatedly selecting adjacent pixels and creates a clean mission trajectory.

4. DQN Model Architecture

The neural network is defined in *model.py*:

```
class DQN(nn.Module):
    def __init__(self, n_actions=100):
        super(DQN, self).__init__()
        self.conv1 = nn.Conv2d(1, 16, kernel_size=3, padding=1)
        self.conv2 = nn.Conv2d(16, 32, kernel_size=3, padding=1)
        self.conv3 = nn.Conv2d(32, 64, kernel_size=3, padding=1)
        self.adaptive_pool = nn.AdaptiveAvgPool2d((4, 4))
        self.fc1 = nn.Linear(64 * 4 * 4, 512)
        self.fc2 = nn.Linear(512, n_actions)
```

Image 6 – Architecture of the DQN network fragment

Why DQN?

- The environment is discrete (every pixel = an action).
- The risk map changes after each action due to the AOE removal.

- Rewards depend on ROI.

Input to the DQN:

A normalized contamination map ($1 \times H \times W$)

Output:

A vector of size $n_actions$ where each entry represents the Q-value of selecting that pixel for testing.

5. Reinforcement Learning Method

The training loop followed standard DQN principles:

- Experience replay buffer
- ϵ -greedy exploration
- Target network
- MSE loss over Q-updates

Reward function:

```
# 2. ROI
roi neto = (probability * VALOR DETECCION) - COSTO TEST
```

Image 7 – ROI calculation fragment

Episodes ended when reward became negative.

6. Frontend Interface (Streamlit)

The Streamlit app performs:


- Parameter tuning (field size, budget)
- Full heatmap visualization
- Path plotting (cyan dashed line)
- Risk matrix coloring (RdYlGn_r)
- Display of mission logs
- ROI metrics

High-density technical grids were created using minor/major locators:

```
ax.grid(which='minor', color='white', linestyle='-', linewidth=0.1, alpha=0.2)
ax.grid(which='major', color='white', linestyle='-', linewidth=0.3, alpha=0.4)
```

Image 8 – Configuration of plot grid and ticks fragment

3. Results



The resulting analysis produced a spatial heatmap of cross-pollination risk over a 1000×500 m soybean field, using a 1 m resolution to approximate real farm-scale conditions. As shown in Figure X, risk values form a clear gradient from the contaminated GMO-adjacent boundary on the right (red, high probability) toward the interior of the non-GMO field (green, low probability), reflecting the exponential distance-decay model and wind-alignment factors defined earlier. Small spatial fluctuations introduced by controlled noise generate realistic “biological texture,” capturing the irregularities in pollen movement caused by micro-turbulence, plant density, and local airflow.

Superimposed on this heatmap are the sampling points selected by the BioSecure AI agent. The system begins by targeting the highest-risk zones at the field edge and progressively moves inward, selecting only those locations where the

expected return on investment (ROI) remains positive. Each blue square marks a sampling action chosen by the agent, while the connecting cyan path shows the mission trajectory. The agent stops automatically once the maximum remaining risk value falls below the economic threshold $p \approx 6.6\%$, at which point additional tests would no longer justify their cost. In this run, the mission achieved a **net ROI of +20,816.58 €**, indicating that targeted sampling prevents unnecessary spending while still covering all meaningful high-risk areas.

4. Discussion and Conclusion

4.1 System Performance and Behavior

The BioSecure AI system demonstrates robust performance across the entire contamination-analysis pipeline, from environmental modeling to reinforcement-learning-based mission planning. The agent consistently identifies high-risk regions and constructs structured inspection paths rather than random or redundant sampling patterns. The inclusion of an ROI-based stopping mechanism adds economic rationality to the behavior: once the expected marginal benefit of additional tests becomes negative, the agent stops autonomously. This behavior mirrors real-world decision processes in farms where testing costs must be balanced against expected contamination penalties.

The system-level results, particularly the generated heatmaps and mission trajectories, show that the hybrid RL + heuristic strategy effectively avoids common pitfalls such as over-exploration of contiguous high-risk zones. Instead, the agent performs distributed sampling across the field, successfully capturing spatial variation in the risk distribution. This validates the effectiveness of the AoE-based cluster-skipping mechanism, which forces the agent to explore new areas instead of repeatedly selecting neighboring high-risk pixels.

4.2 Biological Realism of the Risk Field

A major challenge in agricultural contamination modelling lies in realistically representing pollen dispersion, which is influenced by wind turbulence, crop geometry, and microclimatic fluctuations. This project addresses such complexity through a multi-layered synthetic risk-field generator combining: exponential distance-based decay of pollen probability, structural, spatially coherent noise, and deterministic “debug-mode” matrices.

The structural noise layer significantly improves biological plausibility by breaking unrealistic uniform columns of contamination and introducing natural variability while preserving spatial continuity. This hybrid modelling approach provides the reinforcement-learning agent with a training environment that balances realism and computational efficiency—an essential requirement for real-world deployment.

4.3 Reinforcement Learning Strategy Evaluation

The DQN architecture, coupled with a discrete action space in which each pixel represents an actionable location, is well-suited to the spatial decision-making nature of the task. Reinforcement learning effectively adapts to dynamic states:

after each test, the AoE update reshapes the risk map, making previously high-value regions less informative. This dynamic environment necessitates a learning algorithm that can reason over sequential changes, making DQN an appropriate choice.

Nevertheless, certain RL-specific limitations remain. The action space scales quadratically with field resolution, which increases computational requirements at higher map densities. The hybrid strategy (DQN + AoE + greedy heuristics) mitigates this issue but does not fully solve the scalability challenge. Future work could explore multi-agent RL or attention-based architectures to reduce computational loads.

4.4 Engineering Architecture and Deployment Consideration

The pipeline—consisting of a Flask backend, Streamlit interface, DQN model server, heatmap generator, and mission planner—was designed to be modular, transparent, and deployable. Each component can operate independently:

- 1.The Flask API transforms the system into a service that can integrate with UAVs, mobile devices, or agricultural machinery.
- 2.The Streamlit interface provides immediate usability for non-technical stakeholders such as farmers, inspectors, or certification auditors.
- 3.The mission logs and ROI analytics enhance interpretability and provide a traceable audit trail for regulatory compliance.

The system’s modularity ensures that biological, environmental, or economic models can be replaced or refined without altering the overall architecture. This design makes BioSecure AI suitable not only for research but also for scaling into commercial agricultural applications.

4.5 Limitations and Future Improvements

Although the system achieves promising results, several limitations remain:

(1)Synthetic data dependence

Current experiments rely on simulated risk fields. Real-world contamination dynamics—affected by wind speed fluctuations, temperature gradients, and pollen viability—are more complex than what synthetic fields can fully capture.

(2)Action-space scalability

High-resolution fields introduce very large action spaces. Even with AoE heuristics, training time increases significantly.

(3)Absence of real-field validation

Without field trials, the system cannot yet be certified for regulatory or commercial use.

Future improvements may focus on: integrating UAV-captured imagery and weather-station data; replacing DQN with spatial-attention RL models or multi-agent systems; field-level deployment to measure real contamination, establishing empirical validation.

4.6 Conclusion

This work presents BioSecure AI, an integrated system designed to model, analyze, and mitigate GMO-to-non-GMO cross-pollination risk using a combination of biological simulation, machine learning, and economic optimization. The study demonstrates that reinforcement-learning-driven inspection strategies can significantly reduce unnecessary testing while still covering critical high-risk zones, achieving a balance between cost efficiency and biological safety.

The project makes four key contributions:

1. **A biologically informed contamination model** combining distance-decay functions with structured noise to create realistic and computationally efficient risk maps.
2. **A hybrid RL mission planner (DQN + AoE heuristic)** capable of producing structured, interpretable, and economically rational sampling paths.
3. **An end-to-end deployable pipeline** consisting of a Flask backend, Streamlit interface, heatmap visualizer, and autonomous mission planner.
4. **An ROI-aware decision mechanism** that allows the agent to autonomously determine when further sampling is no longer economically justified.

Together, these elements provide an accessible, low-cost, and scalable tool for farmers, regulatory agencies, and supply-chain operators. The system offers a practical blueprint for integrating AI-driven monitoring into agricultural environments, helping protect the integrity of non-GMO crops and prevent economically damaging contamination events.

Future work will focus on integrating real-world environmental data, expanding the model's predictive capabilities, and conducting validated field trials to bridge the gap between simulation and agricultural deployment. With continued development, BioSecure AI has the potential to become a standard tool for precision agriculture and biological risk management.

5. References

Coléno, F. C., Angevin, F., & Lécroart, B. (2009). A model to evaluate the consequences of GM and non-GM segregation scenarios on GM crop placement in the landscape and cross-pollination risk management. *Agricultural Systems*, 101(1-2), 49-56.

- Ray, J. D., Kilen, T. C., Abel, C. A., & Paris, R. L. (2003). Soybean natural cross-pollination and genetic variability of populations. *Crop Science*, 43(1), 63–70.
- Yoshimura, Y., Beckie, H. J., & Matsuo, K. (2006). Transgenic oilseed rape along transportation routes and port of Vancouver in western Canada. *Environmental Biosafety Research*, 5(2), 67–75.
- Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70–90.
- Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529–533.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- Greene, C. (2016). Economic Issues in the Coexistence of Organic, Genetically Engineered, and Non-GE Crops (USDA-ERS report). U.S. Department of Agriculture, Economic Research Service.
- Brookes, G., & Barfoot, P. (2018). Farm Income and Production Impacts of Using GM Crop Technology, 1996–2016. *GM Crops & Food*, 9, 59–89.
- Yan, S., et al. (2018). Pollen-mediated gene flow from transgenic cotton is constrained by physical isolation measures. *Scientific Reports* (2018).