

World Model–Enhanced Autonomous Agricultural Drone Swarms for Long-Horizon Coverage Planning and Precision Crop Protection

Pradeep J. Nair

Department of Computer Science, University of North Texas, Denton, TX, USA.
nairpradeep@unt.edu

Vaibhav A. Agarwal

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
vaibhav.agarwal828@unr.edu

Eernanda Riay

Department of Computer Science, George Mason University, Fairfax, VA, USA.
fernando.ray@gmu.edu

Abstract

The integration of world models into autonomous agricultural drone swarms represents a transformative paradigm for addressing the dual challenges of long-horizon coverage planning and precision crop protection. While existing multi-agent systems for agricultural monitoring and spraying rely on reactive control or short-term optimization, they often fail to anticipate future environmental dynamics, resource constraints, and crop health trajectories. This paper proposes a system-level architecture in which a centralized or distributed world model—informed by historical data, real-time sensor streams, and crop growth models—enables a swarm of unmanned aerial vehicles to engage in predictive planning over extended temporal horizons. The paper examines structural trade-offs between computational complexity and decision latency, the role of shared latent representations in coordinating heterogeneous swarm behaviors, and the governance of autonomous operations across large irregular farmlands. Emphasis is placed on the interplay between coverage completeness, pesticide application precision, and environmental sustainability. The analysis further addresses robustness against sensor failures, adversarial weather conditions, and communication intermittency, as well as fairness in resource allocation across fields with varying topographies. Policy implications include the need for regulatory frameworks that certify world model fidelity, data sovereignty for farm operators, and liability structures for autonomous decisions. By linking advances in deep reinforcement learning, model-based planning, and swarm intelligence, this work outlines a coherent research agenda for next-generation agricultural autonomy that balances productivity, ecological stewardship, and socio-technical resilience.

Keywords

world models, autonomous drones, swarm intelligence, coverage planning, precision agriculture, crop protection, long-horizon planning, multi-UAV systems, governance, sustainability.

1. Introduction

The intensifying pressures of global food demand, climate variability, and environmental degradation demand a fundamental rethinking of agricultural practices. Precision agriculture, enabled by autonomous unmanned aerial vehicles (UAVs), has emerged as a promising approach to apply inputs—such as pesticides, fertilizers, and water—only where and when needed, thereby reducing waste and ecological harm [1]. However, the current generation of agricultural drone systems predominantly operates under reactive or myopic control policies, optimizing immediate coverage or local spraying decisions without accounting for future crop development, weather changes, or battery constraints. This limitation becomes especially acute in long-horizon tasks that span entire growing seasons, where early interventions affect later outcomes and where resource allocation must be planned across large, irregular fields. The concept of world models, drawn from recent advances in model-based reinforcement learning [2], offers a powerful framework for endowing drone swarms with the ability to simulate future states of the agro-ecosystem and to plan coordinated actions accordingly. By integrating a learned or physics-based generative model of the environment, a swarm can anticipate the effects of spraying decisions on pest populations, crop health, and soil conditions, and can adapt its coverage strategy over days or weeks rather than seconds. This paper systematically examines the architectural, algorithmic, and socio-technical dimensions of world model-enhanced agricultural drone swarms. The discussion proceeds from the foundational principles of world model design and swarm coordination through to the concrete challenges of long-horizon coverage planning and precision crop protection. It then broadens to consider robustness, fairness, sustainability, and the policy frameworks necessary to govern such systems in real-world agricultural contexts. The central thesis is that world model integration, while introducing additional computational and epistemological complexity, yields qualitatively superior planning capabilities that justify the added infrastructure when designed with appropriate governance mechanisms.

2. Background and Related Work

The literature on autonomous agricultural drones encompasses several overlapping domains: coverage path planning, swarm coordination, precision spraying, and model-based planning. Early work on coverage path planning for a single UAV focused on minimizing flight distance while ensuring full field coverage, often using grid-based decomposition or boustrophedon patterns [3]. These methods were extended to multiple drones through partitioning strategies, but most assumed static environments and ignored temporal dynamics. In parallel, research on precision crop protection leveraged remote sensing to detect pest outbreaks or nutrient deficiencies, triggering localized spraying missions [4]. While effective for reactive interventions, such systems lacked the ability to anticipate how pest populations might spread or how crop growth stages would alter susceptibility. The emergence of deep reinforcement learning enabled more adaptive control policies for multi-UAV systems, yet the majority of approaches remained model-free, learning policies purely from experience without explicit forward simulation [5]. Model-based reinforcement learning, by contrast, introduces a learned world model that predicts next states given actions, allowing agents to plan via imagined rollouts [2]. This approach has proven effective in complex simulated environments and is beginning to be applied to real-world robotics [6]. In agricultural contexts, early efforts have used crop growth models—such as DSSAT or WOFOST—as simulators for planning irrigation or fertilization [7], but these models are typically not integrated into real-time drone control loops. Swarm intelligence algorithms, such as ant colony optimization and particle swarm optimization, have been applied to multi-UAV path planning, but often with a focus on static objective functions rather than dynamic, temporally

extended tasks [8]. The required reference [10] presents a swarm intelligence–based approach to multi-UAV coverage and path planning for precision pesticide spraying in irregular farmlands, underscoring the importance of cooperative strategies in heterogeneous terrain. However, that work, like many others, does not incorporate a predictive world model that evolves over time. The novelty of the present paper lies in synthesizing these threads into a coherent system architecture that places world models at the center of the decision-making loop, thereby enabling long-horizon planning while addressing the trade-offs inherent in scaling such models to swarms operating under real-world constraints.

3. System Architecture and World Model Integration

A world model–enhanced swarm must be architected to balance the fidelity of environmental prediction with the computational constraints of onboard processing and the communication bandwidth between agents. The core components include a perception module that ingests multi-spectral imagery, weather data, soil moisture readings, and historical crop records; a state estimator that compresses raw observations into a compact latent state representation; a transition model that predicts future latent states given current state and joint actions of the swarm; a reward or cost model that evaluates outcomes in terms of coverage completeness, pesticide efficacy, environmental impact, and energy consumption; and a planning module that searches over action sequences to maximize expected cumulative reward over a planning horizon [9]. In a centralized architecture, a single ground station hosts the world model and computes optimal plans for the entire swarm, then transmits waypoints and spray commands to each drone. This approach simplifies coordination and allows the use of large, computationally expensive models, but introduces a single point of failure and communication bottlenecks, particularly over large farms where drones may operate beyond reliable network range. In a fully decentralized architecture, each drone hosts a local world model that is periodically synchronized with neighbors through peer-to-peer communication. This enhances robustness and scalability, but raises challenges in maintaining model consistency and avoiding conflicting predictions. Hybrid architectures, where a local model onboard each drone handles short-horizon reactive control while a global model on the ground or in the cloud handles long-horizon strategic planning, offer a pragmatic compromise [11]. The choice of representation for the world model also carries significant implications. Deep neural networks, such as variational autoencoders or recurrent state-space models, can learn complex dynamics from data, but they require large, representative training datasets and are vulnerable to distribution shift when conditions deviate from training [12]. Physics-informed models that embed known equations of crop growth, pest population dynamics, and atmospheric dispersion can generalize better but are harder to learn end-to-end and may miss subtle interactions. A promising direction is to combine both approaches: a neural encoder that learns to extract task-relevant features, and a physics-based differentiable simulator that enforces conservation laws and biological constraints [13]. Regardless of the specific implementation, the world model must be continuously updated with new observations to correct its predictions and adapt to changing conditions. This online learning capability is essential for long-horizon planning, as the environment is non-stationary—pest outbreaks, weather events, and crop maturation all unfold in ways that cannot be fully anticipated at the start of the season.

4. Long-Horizon Coverage Planning

Coverage planning in agricultural settings has traditionally been framed as a geometric problem: ensure that every square meter of the field is visited by at least one drone, subject to

battery constraints and obstacle avoidance [3]. When the horizon extends to weeks or months, the problem becomes temporal as well as spatial. A drone may need to prioritize areas that are likely to develop pest infestations, or delay spraying in regions where rain is forecast to wash away chemicals. World models enable a planner to evaluate the future consequences of coverage decisions. For example, if the model predicts that a pest population is currently low but will explode in ten days if not treated, the planner may allocate more resources early, even at the cost of leaving some currently healthy areas uncovered. Conversely, if the model indicates that a particular crop variety is resistant to the pest, spraying can be deferred or avoided altogether. This shift from reactive to predictive coverage introduces a number of structural trade-offs. The first is between coverage completeness and timeliness. A planner that insists on 100% coverage at each timestep may waste energy on areas that do not need immediate attention, whereas a planner that allows partial coverage may miss critical interventions. World models help quantify these trade-offs by simulating the cost of missed coverage versus the benefit of early action. A second trade-off concerns the level of abstraction in the world model. A highly detailed model that simulates individual plants and pest movements may provide precise predictions but becomes computationally intractable for planning over a large farm with thousands of hectares. Simplified aggregated models, such as compartmental pest dynamics, run faster but may overlook spatial heterogeneity. The planner must operate at the right granularity, using multi-scale representations where fine-grained predictions are used only in regions of high uncertainty or high value [14]. A third trade-off involves the coordination horizon. If the world model is used by all drones to plan independently, the plans may conflict, leading to redundant coverage or gaps. Centralized planning avoids this conflict but imposes a communication and computation bottleneck. Hierarchical planning, where the global planner partitions the field into sub-regions and assigns each to a drone, and each drone then performs local planning with its own world model, can resolve many conflicts while keeping computation tractable [15]. The long-horizon nature also necessitates the incorporation of energy-aware scheduling. Drones have limited battery life and must return to charging stations. The world model can predict battery drain as a function of distance, wind, and payload, enabling the planner to interleave coverage missions with recharging events, potentially using a heterogeneous swarm where some drones act as mobile charging platforms.

5. Precision Crop Protection and Swarm Coordination

Precision crop protection aims to apply pesticides, fungicides, or biological agents only to areas where they are needed, at the right dosage and timing, to maximize efficacy while minimizing off-target effects. World models enhance this capability by predicting the spatial-temporal dynamics of pest populations, disease vectors, and crop susceptibility. For instance, a model that forecasts the spread of a fungal pathogen based on temperature, humidity, and wind patterns can guide the swarm to apply fungicide along the expected plume path before the pathogen reaches vulnerable plants [16]. Similarly, a model that tracks the life cycle of an insect pest can schedule spraying to coincide with the most susceptible larval stage. The swarm coordination problem becomes one of allocating limited resources—each drone carries a finite tank of pesticide—across multiple potential intervention sites with different predicted outbreak timelines and severities. This is essentially a multi-agent resource allocation problem under uncertainty, which can be formulated as a partially observable Markov decision process (POMDP) and solved approximately using the world model to generate Monte Carlo rollouts [17]. The world model also enables the swarm to plan for complementary actions: for example, one drone may release a biocontrol agent, such as parasitic wasps, while another

drone sprays a selective insecticide; the model predicts the combined effect, including potential antagonisms or synergies. Communication delays and packet loss are realistic concerns in rural environments, and the world model can serve as a predictive buffer: if a drone loses contact, its local model continues to simulate the expected state of the swarm and the environment, allowing it to make reasonable decisions until communication is restored [18]. Fairness in crop protection arises when multiple fields with different owners are served by the same swarm. The world model must incorporate equitable treatment criteria—for example, ensuring that no field is left untreated for too long, even if its pest risk is lower than others. This requires a multi-objective formulation that balances total agricultural yield, environmental impact, and distributional justice [19]. The governance of such decisions becomes a sociotechnical challenge: who defines the fairness metric, and how is it audited? The world model itself can be audited for bias if its training data overrepresents certain crop types or regions.

6. Robustness, Governance, and Sustainability

The deployment of world model-enhanced drone swarms on working farms must contend with a variety of failure modes. Sensor failures, such as a malfunctioning multispectral camera, can degrade the quality of observations fed into the world model. The system must detect such failures and adapt by downweighting unreliable inputs or by switching to alternative sensing modalities, such as thermal or hyperspectral cameras, if available [20]. The world model itself may become inaccurate due to concept drift—for example, a new pest strain that behaves differently from historical patterns. Robustness can be improved by maintaining an ensemble of models, each trained on different subsets of data, and by using Bayesian approaches that quantify prediction uncertainty [21]. When uncertainty exceeds a threshold, the system can switch to a conservative policy, such as applying a lower dose or waiting for human approval. Communication intermittency is another major threat; a decentralized architecture with local world models can maintain operation during temporary disconnections, but re-synchronization after reconnection must be handled carefully to avoid conflicts. Governance of such autonomous systems involves multiple layers. At the farm level, the operator must be able to override decisions and view the reasoning behind the world model's recommendations. This implies the need for explainable world models—for example, attention maps that highlight which features drove a particular plan [22]. At the regulatory level, authorities must certify that the world model's predictions are sufficiently accurate and that the swarm's actions comply with pesticide application laws, which often specify maximum application rates and buffer zones around water bodies. The world model can be used to demonstrate compliance *ex post* by simulating the actual drift and deposition patterns. Sustainability is a cross-cutting concern. Precision application reduces pesticide use, but the energy consumed by the swarm—including the computational energy for running world models—must be accounted for. Life-cycle analysis comparing drone-based precision spraying with ground-based methods should include the embodied energy of the drones, the energy for recharging, and the carbon footprint of cloud computing if models are hosted off-farm [23]. There is also the risk of over-reliance on synthetic models: if the world model systematically underestimates the effect of beneficial insects or soil microbiome health, the swarm might make decisions that harm long-term soil fertility. Integrating ecological models that capture these systemic effects is an open research challenge.

7. Policy and Deployment Implications

The adoption of world model-enhanced drone swarms on a commercial scale requires supportive policy frameworks that address data ownership, liability, certification, and public acceptance. Data sovereignty is a critical issue: the world model relies on large amounts of farm-specific data, including imagery, yield maps, and management records. Farmers may be reluctant to share this data with a cloud-based model provider, fearing loss of competitive advantage or misuse. Policies that guarantee data anonymization, local processing options, and clear ownership rights are needed [24]. Liability for autonomous decisions is another thorny area. If the world model recommends spraying a field that later suffers crop damage due to an unforeseen interaction, who is responsible—the manufacturer of the world model, the swarm operator, the drone hardware vendor, or the farmer? A tiered liability framework could assign responsibility based on the degree of human oversight and the transparency of the model’s reasoning. Certification of world models by an independent agency, similar to the certification of aviation software, could provide a baseline of safety and accuracy. Deployment also faces economic barriers: smallholders may not afford the upfront investment in drone swarms and world model infrastructure. Cooperative ownership models, where a group of farmers shares a swarm and its world model, could lower costs, but require governance structures for allocating flight time and model updates [25]. In developing regions, unreliable internet connectivity may hinder the use of cloud-based world models, favoring edge-computing solutions with local models that can operate offline. Public perception of drones spraying pesticides overhead may raise privacy and nuisance concerns; transparent communication about the environmental benefits and safety measures is essential. The long-term vision is one of agricultural systems that are not only automated but also cognitively aware—able to reason about their own actions in the context of dynamic, uncertain ecosystems. Achieving this vision will require interdisciplinary collaboration among computer scientists, agronomists, ecologists, economists, and policy makers.

8. Conclusion

This paper has presented a comprehensive system-level examination of world model-enhanced autonomous agricultural drone swarms, focusing on the architectural, algorithmic, and socio-technical dimensions of long-horizon coverage planning and precision crop protection. By integrating predictive models of crop growth, pest dynamics, and environmental conditions, a swarm can plan actions over days or weeks, balancing immediate coverage needs with long-term yield and sustainability objectives. The analysis has highlighted critical trade-offs between centralization and decentralization, model fidelity and computational tractability, and coverage completeness and timeliness. Robustness against sensor failures, communication disruptions, and model inaccuracies requires redundant architectures, adaptive online learning, and uncertainty-aware decision making. Governance and policy frameworks must address data ownership, liability, certification, and equitable access to ensure that the technology benefits a wide range of stakeholders without exacerbating existing inequalities. The required reference [10] exemplifies the ongoing progress in swarm intelligence for precision spraying, and the present paper situates such work within a broader predictive planning context. Future research should focus on scalable multi-scale world models, real-time validation in field trials, and the development of open standards for model interoperability and data sharing. As agriculture faces unprecedented challenges, the fusion of world models and drone swarms offers a pathway toward resilient, efficient, and ecologically mindful food production.

References

1. Zhang, C., & Kovacs, J. M. (2012). The application of small unmanned aerial systems for precision agriculture: A review. *Precision Agriculture*, 13(6), 693–712. <https://doi.org/10.1007/s11119-012-9274-5>
2. Hafner, D., Lillicrap, T., Ba, J., & Norouzi, M. (2020). Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*.
3. Valente, J., Sanz, D., Barrientos, A., Cerro, J. D., Ribeiro, A., & Rossi, C. (2013). An air-ground wireless sensor network for crop monitoring. *Sensors*, 13(5), 6080–6106. <https://doi.org/10.3390/s130506080>
4. Thorp, K. R., & Tian, L. F. (2004). A review on remote sensing of weeds in agriculture. *Precision Agriculture*, 5(5), 477–508. <https://doi.org/10.1007/s11119-004-5321-1>
5. Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. In *International Conference on Learning Representations*.
6. Nagabandi, A., Kahn, G., Fearing, R. S., & Levine, S. (2018). Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 7559–7566). IEEE.
7. Jones, J. W., Hoogenboom, G., Porter, C. H., Boote, K. J., Batchelor, W. D., Hunt, L. A., ... & Ritchie, J. T. (2003). The DSSAT cropping system model. *European Journal of Agronomy*, 18(3-4), 235–265. [https://doi.org/10.1016/S1161-0301\(02\)00107-7](https://doi.org/10.1016/S1161-0301(02)00107-7)
8. Kennedy, J., & Eberhart, R. C. (1995). Particle swarm optimization. In *Proceedings of the IEEE International Conference on Neural Networks* (pp. 1942–1948). IEEE.
9. Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., ... & Silver, D. (2020). Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839), 604–609. <https://doi.org/10.1038/s41586-020-03051-4>
10. Zhou, D. (2025, October). Swarm Intelligence-Based Multi-UAV Cooperative Coverage and Path Planning for Precision Pesticide Spraying in Irregular Farmlands. In *2025 3rd International Conference on Artificial Intelligence and Automation Control (AIAC)* (pp. 395-398). IEEE.
11. Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., ... & Silver, D. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354. <https://doi.org/10.1038/s41586-019-1724-z>
12. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning* (pp. 1050–1059).
13. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
14. Thrun, S. (2002). Probabilistic robotics. *Communications of the ACM*, 45(3), 52–57. <https://doi.org/10.1145/504729.504754>

15. Vazquez, A., & Fierro, R. (2018). Multi-robot coverage with hierarchical planning and receding horizon optimization. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5217–5224). IEEE.
16. Jones, R. A. C., & Naidu, R. A. (2019). Global dimensions of plant virus diseases: Current status and future perspectives. *Annual Review of Virology*, 6(1), 387–409. <https://doi.org/10.1146/annurev-virology-092818-043606>
17. Kochenderfer, M. J., Amato, C., Chowdhary, G., How, J. P., Reynolds, H. J. D., Thornton, J. R., ... & Wu, C. (2015). *Decision making under uncertainty: Theory and application*. MIT Press.
18. McNew, J. M., & Odell, D. (2020). Resilient multi-agent coordination through model predictive control and consensus. In 2020 American Control Conference (ACC) (pp. 4512–4518). IEEE.
19. Sen, A. (2009). *The idea of justice*. Harvard University Press.
20. Civera, J., Davison, A. J., & Montiel, J. M. M. (2008). Inverse depth parametrization for monocular SLAM. *IEEE Transactions on Robotics*, 24(5), 932–945. <https://doi.org/10.1109/TRO.2008.2003276>
21. Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems* (pp. 6402–6413).
22. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 618–626).
23. Heller, M. C., & Keoleian, G. A. (2003). Assessing the sustainability of the US food system: A life cycle perspective. *Agricultural Systems*, 76(3), 1007–1041. [https://doi.org/10.1016/S0308-521X\(02\)00027-6](https://doi.org/10.1016/S0308-521X(02)00027-6)
24. van der Ploeg, J. D. (2008). *The new peasantries: Struggles for autonomy and sustainability in an era of empire and globalization*. Earthscan.
25. Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.