Volume 4, Issue 6, June - 2026
ISSN (E): 2938-3781

Web of Agriculture: Journal of Agriculture and Biological Sciences
webofjournals.com/index.php/8

FORECASTING HUNGER BEFORE THE HARVEST: A SOVEREIGN, EARLY, AND UNCERTAINTY AWARE CROP YIELD AND FOOD SECURITY EARLY WARNING FRAMEWORK FOR CLIMATE EXPOSED IRRIGATED SYSTEMS, WITH AN ORIGINAL HYBRID OBSERVATION AND SIMULATION ENGINE (HOSIL)

Primov Abdulla Egamkulovich
Vice Rector for Scientific Research,
Innovation and Spirituality,
International Agriculture University
E-mail: Abdulla.Primov@rau.ac.uk
ORCID: 0000-0002-1889-2161

Ismoilov Avazbek Abdusamat o‘g‘li
BSc Agrologistics, 2nd year student
International Agriculture University
E-mail: avazbeki733@gmail.com
ORCID: 0009-0009-9659-1514

Abstract

In 2025 the world received an unwelcome demonstration of how fragile its hunger warning infrastructure had become. The Famine Early Warning Systems Network, the gold standard that had for four decades forecast food crises months in advance, was abruptly taken offline in January when its single donor was dismantled, and although it resumed later in the year, the episode exposed a structural dependence that no food insecure country should accept. This paper argues that the appropriate response is not merely to restore the old arrangement but to build sovereign, low cost, scientifically rigorous national capacity for crop yield and food security forecasting that can survive donor shocks, that warns early enough to enable action rather than only to record disaster, and that is honest about its own uncertainty. We synthesize the established discipline of agricultural early warning and crop yield forecasting with the current methodological frontier, namely the fusion of process-based crop simulation models, machine learning, and Earth observation data assimilation, and with the operational practice of anticipatory action, in which pre-arranged financing is released automatically when forecast triggers are met. As the central contribution we propose an original system, not yet implemented, which we name HOSIL, a Hybrid Observation and Simulation Intelligence for Livelihoods. HOSIL assimilates Earth observation state variables derived from geospatial foundation model embeddings and multitemporal radar and optical data into a calibrated crop model, couples that mechanistic core with a label efficient machine learning correction trained partly on synthetic data, forecasts yield progressively through the season with shrinking and quantified uncertainty, translates the forecast distribution into a national food balance, and converts that balance into costed anticipatory action triggers. We specify the architecture and the rationale for each design choice, set out a validation design that could refute the approach, and analyze its failure modes without flinching, including the cost of false alarms, the blindness of crop models to conflict driven hunger, the capacity required for sovereign operation, and the moral hazard of automated triggers. We close with prioritized recommendations and a phased, buildable roadmap. The contribution is a defensible architecture for turning open data and open models into timely, accountable, and nationally owned warnings that are followed by action.

Keywords: Food security; famine early warning; crop yield forecasting; data assimilation; hybrid crop modeling; geospatial foundation models; anticipatory action; forecast based financing; uncertainty; Central Asia; winter wheat; sovereign monitoring.

Introduction

A warning that arrives too late, or that nobody owns, or that nobody acts on, is not a warning. This unglamorous truth sits at the center of the agricultural early warning enterprise, and it was thrown into sharp relief in 2025. For four decades the Famine Early Warning Systems Network had monitored weather, crop production, prices, and conflict across more than thirty of the most food insecure countries, producing projections of acute food insecurity six to nine months ahead and serving as the reference that humanitarian agencies and national governments relied upon. In January 2025 that system was abruptly taken offline when its sole funder was dismantled, and although it resumed operations later in the year, the interruption left a gap precisely when conditions were deteriorating, and it revealed a dependence on a single donor that is incompatible with the seriousness of the task. The companion framework for classifying the severity of food crises was simultaneously strained. The lesson is not that early warning failed but that its ownership was concentrated in a way that made it brittle.

This paper takes that lesson as its starting point and asks a constructive question. Given the open satellite archives, the open foundation models, and the open crop models now available, can a food insecure country or region build its own crop yield and food security forecasting capacity that is scientifically rigorous, early, honest about uncertainty, and resilient to the withdrawal of any single external donor? We argue that it can, that the scientific ingredients are now mature, and that the binding constraints are integration, calibration, and the link to action rather than the availability of data or methods.

Three bodies of knowledge must be brought together to answer the question. The first is the long and disciplined tradition of agricultural early warning and crop yield forecasting, which established that forecasting is the foundation of any warning system, that crop simulation models translate weather, soil, and management into yield, that a forecast without a probability rating is of limited use, and that the more distant the forecast the less reliable it becomes, so that timeliness and reliability must be traded against each other deliberately. National and international programmes such as the agrometeorological forecasting services of several countries, the staged preharvest forecasts issued at multiple points in the crop cycle by national crop forecast centers, the agricultural stress indices maintained by international agencies, and the coordinated global monitoring of the major breadbasket crops, all embody this tradition and all demonstrate that

Web of Agriculture: Journal of Agriculture and Biological Sciences
webofjournals.com/index.php/8

remote sensing has moved from research into routine production for crop area and increasingly for yield.

The second body of knowledge is the methodological frontier in yield prediction. The field has spent a decade discovering that the two dominant approaches, process based crop models that are difficult to calibrate locally and data driven machine learning that is hungry for labels, are complementary rather than opposed. The current state of the art couples them: Earth observation state variables such as leaf area index and soil moisture are assimilated into a crop model, the data assimilated model is then corrected by machine learning that captures effects the mechanistic model omits, and in the most label scarce settings the crop model is even used to generate synthetic data to pretrain a network that is then finetuned on the few real observations available. In parallel, geospatial foundation models now supply analysis ready, cloud robust, label efficient representations of the land surface, including at ten meter resolution, that can serve as the observation backbone for exactly this kind of assimilation.

The third body of knowledge is the operational practice that closes the loop between warning and action. Anticipatory action, originally called forecast based financing, pre arranges the amount, source, and disbursement of funds so that a defined response can begin automatically the moment a forecast trigger is crossed, before the hazard impacts unfold. Over a decade this practice has grown from pilots into a portfolio spanning more than forty countries, reaching millions of people with pre arranged finance, and it has shifted the humanitarian system from reacting to hazards toward acting ahead of them. Its weakest dependency, however, is the quality and ownership of the forecast that pulls the trigger.

Placing these three together reveals the opportunity. The forecasting tradition tells us what a warning must contain and that it must be probabilistic and timely. The methodological frontier tells us how to produce yield forecasts that are accurate and label efficient even where ground data are scarce. Anticipatory action tells us how to convert a forecast into protection for real people. Yet no widely available system unites all three under national ownership in a form that is resilient to donor shocks. That is the gap this paper sets out to fill, and the central contribution is an original, not yet implemented system, named HOSIL, designed precisely to occupy it.

We make four contributions. First, we frame the 2025 disruption as a structural argument for sovereign forecasting capacity rather than as a transient funding accident. Second, we synthesize the early warning tradition, the hybrid crop model and machine learning frontier, and anticipatory action into a single coherent design. Third, we specify the HOSIL architecture component by component, with the rationale for each choice and the alternatives rejected, and we set out a validation design that could prove it wrong. Fourth, we are candid about the failure modes, because a forecasting system that pulls financial triggers and that may justify or withhold assistance to vulnerable people must be judged by its calibration and its humility, not only by its accuracy.

Background and state of the art

The early warning and crop yield forecasting tradition. Agricultural early warning rests on a few durable principles. Forecasting is the foundation of all warning, and it must address the availability, stability, access, and utilization dimensions of food security with enough lead time for decision makers to react and with enough reliability to avoid the false alarms that erode trust. Every forecast carries a probability rating, and as a general rule longer horizons buy lead time at the cost of reliability, so the design of any system is in part a deliberate choice of where on that tradeoff to operate. Crop simulation models, which represent crop growth and yield as functions of soil, weather, and management, are the workhorse of yield forecasting, ranging from simple statistical relationships used in national early warning systems to complex mechanistic representations, and they are known to omit important real world factors such as weeds, pests, diseases, and nutrient limitation. Operational programmes demonstrate the maturity of the field. National crop forecast centers issue a sequence of preharvest forecasts at successive stages of the season, blending agrometeorological models, medium resolution remote sensing through the middle of the season, and high resolution imagery near harvest, and they sustain accuracy by conducting field based crop cutting experiments to calibrate their models each season. International agencies maintain agricultural stress indices built on dekadal vegetation and temperature data to flag the share of cropland under drought, and coordinated global monitoring tracks the condition of the major traded crops across producing regions.

The Earth observation yield chain. Remote sensing contributes to this tradition at every link of a chain that runs from imagery to decision. Anomaly detection and vegetation condition flag where conditions depart from normal; crop type and area mapping establish what is planted where; crop state variables such as leaf area index and biomass feed growth models; yield is estimated and aggregated to production; and production combined with stocks, demand, and prices yields a food balance and a food security assessment. Different links demand different sensor characteristics, from frequent coarse observation for condition monitoring to fine resolution near harvest for district level estimates. Foundation models now strengthen the early links of this chain by supplying transferable, gap robust representations that reduce the label burden of crop mapping and state variable retrieval, which is decisive where field campaigns are scarce.

The hybrid frontier: process models, machine learning, and data assimilation. The most consequential recent development for yield forecasting is the reconciliation of mechanistic and data driven modeling. Process based crop models encode agronomic knowledge but are hard to calibrate to local conditions; machine learning fits local data well but needs a great deal of it and extrapolates poorly. Hybrid approaches combine their strengths. One family assimilates Earth observation state variables, such as leaf area index and soil moisture, into the crop model so that the simulation is continually corrected by what the satellite actually sees, then applies machine learning to capture residual effects that the mechanistic model does not represent, with documented gains for cereal yield forecasting. A second family uses the crop model as a generator of synthetic training data to pretrain a network, which is then finetuned on the limited real observations, a meta modeling strategy that improves predictions precisely in the data scarce regime that characterizes most food insecure regions. These advances matter here because they make accurate yield forecasting feasible without the dense ground truth that data poor countries cannot supply.

From early warning to early action: anticipatory action. A forecast protects no one until it triggers a response. Anticipatory action institutionalizes the link by agreeing in advance on three elements: the trigger, usually a forecast threshold; the actions to be taken before the hazard; and the amount and source of pre arranged finance, so that funds disburse immediately when the trigger is crossed. The approach has matured from early pilots into a portfolio across more than forty countries, reaching millions of people, and its evidence base shows that acting ahead of a shock reduces humanitarian need relative to reacting after it. The persistent vulnerability of the model is the

forecast itself: a poorly calibrated trigger either wastes scarce funds on false alarms or fails to fire when it should, which places the scientific quality and the institutional ownership of the forecast at the very center of whether anticipatory action succeeds.
The unoccupied position. Each of these strands is strong on its own, and yet the combination that the moment demands does not exist as an accessible, nationally owned system. The forecasting tradition is mature but its flagship implementations have proven donor dependent. The hybrid methodological frontier is advancing rapidly but largely in research settings and breadbasket geographies rather than in operational, sovereign, food insecure contexts. Anticipatory action is scaling but is bottlenecked by forecast quality and ownership. The unoccupied position is a sovereign, low cost, early, explainable, uncertainty aware yield and food security forecasting system whose outputs are designed from the start to drive anticipatory action and to survive the withdrawal of any single external funder. HOSIL is our proposal to occupy it.

HOSIL: an original proposed system (not yet implemented)

This section presents the author original design contribution. HOSIL, a Hybrid Observation and Simulation Intelligence for Livelihoods, is a proposed architecture rather than a completed system, and it is described here at the level of design and rationale so that it can be built, tested, and criticized. The name is chosen deliberately, since hosil is the local word for harvest, and the system exists to forecast that harvest early enough and credibly enough to protect the people who depend on it.

Design principles. Five principles govern every subsequent choice. The system must be sovereign, meaning it runs on open satellite archives, open foundation models, and open crop models so that a national agency can operate it without dependence on a single external donor, which is the direct structural response to the fragility exposed in 2025. It must be early, meaning it produces useful forecasts months before harvest rather than confirmations after it, because lead time is what makes action possible. It must be explainable, because a water or food authority will not and should not act on an opaque number, so every forecast is accompanied by the physical quantities behind it, the crop condition, the simulated growth, and the assimilated observations. It must be uncertainty aware, meaning it outputs a forecast distribution with a calibrated probability rather than a point estimate, honoring the founding principle that a forecast without a reliability rating is of limited use. And it must be action linked, meaning its outputs are expressed in the units that anticipatory action requires, namely triggers, expected impact, and costed response, rather than in units convenient only to the modeler.

The observation layer. The observation layer turns raw satellite streams into the state variables that the crop model needs and into the spatial backbone for crop area. It draws crop condition and state variables, in particular leaf area index, biomass proxies, evapotranspiration, and soil moisture, from multitemporal radar and optical observation, and it uses geospatial foundation model embeddings as a cloud robust, label efficient backbone for crop type and area mapping. The rationale is that the assimilation core is only as good as the observations it ingests, and foundation model representations reduce both the label burden and the sensitivity to clouds and missing data that otherwise corrupt state variable time series in the critical weeks of the growing season.

The hybrid simulation and machine learning core. The core is deliberately hybrid. A process based crop model simulates the growth and yield of the target crop from weather, soil, and management, providing agronomic structure and the ability to reason about conditions never seen in the training data. Earth observation state variables are sequentially assimilated into this model so that the simulation is continually corrected toward what the satellite observes, rather than drifting on the strength of its assumptions. A machine learning component then corrects the residual error of the data assimilated simulation, capturing the pests, diseases, management heterogeneity, and extreme events that mechanistic models omit. Two design decisions follow from the label scarce context. The machine learning component is kept lightweight and is pretrained on synthetic data generated by the crop model itself, then finetuned on the few real yield observations available, which is the meta modeling strategy that has been shown to outperform purely data driven baselines when data are scarce. And the mechanistic core guarantees a physically plausible forecast even where the data driven correction has little local support, which protects against the confident nonsense that ungrounded models can produce out of distribution.

Progressive in season forecasting. HOSIL does not wait for harvest. It issues a sequence of forecasts at successive points in the season, beginning early with wide uncertainty driven mainly by the crop model and seasonal weather expectations, and tightening as the season advances and more observations are assimilated, in the spirit of the staged preharvest forecasts that mature national centers already issue. The value of this design is that it places the lead time and reliability tradeoff in the hands of the decision maker rather than hiding it: an early forecast with honestly wide intervals can justify low regret preparatory actions, while a later forecast with narrow intervals can justify costly committed ones. Crucially, the uncertainty is quantified and decreasing in a way the user can see, which is what allows a trigger to be set at a defensible point on the curve.

From yield distribution to a national food balance. A district yield forecast is not yet a food security statement. HOSIL aggregates the spatially explicit, uncertainty bearing yield forecasts to subnational and national production, combines them with planted area from the mapping backbone, and sets the result against consumption requirements, stocks, trade, and price signals to produce a probabilistic national food balance. This is the step that converts an agronomic quantity into the availability and stability dimensions of food security, and expressing it as a distribution rather than a single number is what makes it usable for risk based decisions. Where the system cannot observe a driver, most importantly conflict and market disruption, it says so explicitly rather than implying a completeness it does not have.

The anticipatory action engine. The final layer is what distinguishes HOSIL from a forecasting exercise. It converts the probabilistic food balance into the three elements that anticipatory action requires. It defines triggers as thresholds on the forecast distribution, for example a stated probability that production in a region will fall below a critical level by a stated date. It links each trigger to a menu of costed actions appropriate to the lead time, from low regret preparatory measures at long lead to committed transfers at short lead. And it is designed to interface with pre arranged finance so that, subject to human authorization, crossing a trigger can release funds before the shock matures. The engine is explicitly calibrated to manage the cost of false alarms against the cost of missed events, because in anticipatory action a trigger set too sensitively wastes scarce funds and erodes trust while a trigger set too conservatively fails the people it was meant to protect.

Why a coupled, sovereign design rather than a chain of tools. HOSIL could in principle be assembled from independent components, but three properties require co design. Uncertainty must propagate coherently from observation error through assimilation, simulation, the machine learning correction, aggregation, and the food balance, so that the trigger reflects all sources of error rather than a single overconfident point. Explainability must be preserved end to end, so that a trigger can be traced back to the crop condition and the assimilated observations that produced it. And sovereignty must be engineered, not assumed, by selecting at every layer an open and substitutable component so that the failure or withdrawal of any one provider does not disable the system, which is the lesson of 2025 translated into architecture.

Proposed implementation

Demonstration setting. A natural first demonstration is winter wheat in the irrigated systems of Central Asia, a crop and region where the cereal harvest is central to national food security, where the growing calendar is well defined, where water and heat stress create real interannual variability, and where dependence on external monitoring is a strategic vulnerability. Winter wheat is well suited to progressive forecasting because its long season offers multiple assimilation points between emergence and harvest, and the regional history of cereal yield forecasting from satellite data provides a foundation to build upon.

Data and openness. The system is specified on open inputs by design: open optical and radar archives and harmonized products for observation, open geospatial foundation model embeddings for the mapping and state variable backbone, open meteorological reanalysis and seasonal forecasts for the crop model drivers, open soil databases, and an open crop model. Reference yields come from official statistics and, where feasible, a modest program of crop cutting experiments to calibrate and validate, following the practice of mature national forecast centers. The deliberate consequence of this choice is that the running system has no single point of external dependence.

Label efficient training and calibration. Because local yield labels are scarce, training leans on three devices already justified above: synthetic pretraining from the crop model, finetuning on the few real observations, and the mechanistic core as a physically grounded prior. Calibration is treated as a first class task, not an afterthought, because the entire value of the anticipatory action engine depends on the forecast probabilities meaning what they say. The system is calibrated and validated against historical seasons, with particular attention to the dry and shock years that the triggers most need to catch.

Validation design

A proposal of this kind earns belief only through a validation that could refute it. We specify the following.

Forecast skill as a function of lead time. The primary evaluation is the skill of the yield forecast at each lead time, reported as error against official statistics and against any crop cutting experiments, and crucially as skill relative to a naive climatological baseline, because beating climatology at long lead is the real test of value. The expected and desired result is a skill curve that rises as the season advances, and the honest reporting of how early the system becomes useful is itself a key finding.

Calibration and probabilistic quality. Because the outputs are distributions that drive triggers, calibration is evaluated directly with reliability diagrams and proper scoring rules, testing whether events assigned a given probability occur at that frequency. A system that is accurate on average but miscalibrated is dangerous in an anticipatory action setting and must be reported as such.

Food balance and trigger performance. The national food balance is validated against subsequent official production and food security assessments. The anticipatory action engine is evaluated by hindcasting historical shock years to ask whether its triggers would have fired, with what lead time, and at what rate of false alarms and missed events, since these operational quantities, not abstract accuracy, determine whether the system would protect or fail people in practice.

Baselines, ablations, and transfer. The baselines are strong: a pure crop model without assimilation, a pure machine learning model, and a simple vegetation index regression of the kind used in many operational early warning systems. Ablations remove the data assimilation, the machine learning correction, and the synthetic pretraining in turn to isolate the contribution of each. Transfer is tested across regions and across held out years, including years not seen in training, because a system intended for sovereign operation across a heterogeneous region must degrade gracefully rather than collapse outside its training distribution.
Success criteria stated in advance. To prevent post hoc rationalization, success is defined before results are seen. HOSIL succeeds if it beats the climatological and vegetation index baselines on yield skill at a policy relevant lead time of at least several months; if its forecast probabilities are well calibrated; if its triggers would have fired with useful lead time on historical shock years at an acceptable false alarm rate; and if it retains usable skill under cross region and cross year transfer. Failing any criterion is a result to publish, not to hide.

Anticipated results, tradeoffs, and honest uncertainty
We state expectations as reasoned hypotheses with attached confidence, not as findings.
We expect, with moderate to high confidence, that the hybrid core will outperform both the pure crop model and the pure machine learning baseline, because the complementarity of mechanistic structure and data driven correction is now well documented across cereal systems. We expect, with moderate confidence, that synthetic pretraining will materially reduce the real labels needed for a given skill, which is the property that makes sovereign deployment realistic in data scarce settings. We expect, with lower confidence, that the system will achieve calibrated and useful skill at the several month lead time that anticipatory action prefers, because skill at long lead is genuinely hard and depends heavily on seasonal weather predictability, which is the least controllable input. We are frankly uncertain whether the trigger performance will reach an operationally acceptable balance of false alarms and missed events, and we regard establishing

that balance as the decisive empirical question rather than a detail.
The tradeoffs are real. Earliness buys lead time at the cost of reliability, the founding tension of all forecasting, and HOSIL addresses it not by pretending to resolve it but by exposing it as a visible, decreasing uncertainty that the decision maker can act against. Sovereignty buys resilience at the cost of demanding national technical capacity, since a system nobody can run is no more sovereign than one a donor switches off. Explainability buys trust at some cost in raw accuracy, since the most accurate black box may be less useful than a slightly less accurate model whose triggers can be audited. And probabilistic honesty buys good decisions at the cost of communicative simplicity, since a distribution is harder to act on than a number, which places real weight on the design of the trigger and the interface.

Risks, failure modes, and limitations

This section is adversarial toward the proposal on purpose.The false alarm problem is fundamental and not merely technical. Severe production shocks are rare, so even a good forecast crossing a trigger will often be a false alarm in absolute terms, and repeated false alarms exhaust scarce funds and destroy the credibility on which anticipatory action depends. The honest response is to design triggers around the cost of action versus the cost of inaction, to prefer low regret actions at long lead, and to report the base rate explicitly rather than quoting accuracy in a vacuum.
Crop models are blind to the largest drivers of modern food crises. Conflict, displacement, market collapse, and policy shocks cause much of the world acute hunger, and no amount of satellite assimilated agronomy will see them coming. HOSIL therefore forecasts the agricultural availability dimension well and must be explicit that it is only one input into a food security assessment that requires market, conflict, and access information from other sources. Presenting an agronomic forecast as a complete food security forecast would be a dangerous overreach.
Sovereignty is a capacity claim as much as a technical one. Open data and open models remove the donor dependence of inputs, but a national agency still needs the people, computation, and institutional continuity to run the system, and without sustained investment in that capacity the sovereignty is nominal. The 2025 episode is a warning against single points of failure of every kind, including a single overstretched national team.
Transfer and nonstationarity threaten reliability. A model calibrated on past seasons may mislead under a changing climate that pushes conditions outside the historical envelope, which is exactly when warnings matter most. The mechanistic core mitigates this by reasoning physically about unseen conditions, but it does not eliminate the risk, and continual recalibration is required.
Finally, there is a moral hazard in automated triggers and an equity hazard in acting on uncertain forecasts. A trigger that releases or withholds assistance based on a model can entrench errors at scale and can be gamed once its rules are known. The appropriate posture is to keep human authorization in the loop, to treat the forecast as decision support rather than decision, to surface uncertainty rather than bury it, and to design so that the cost of model error does not fall on those least able to bear it.

Policy and governance implications

The implications follow directly from the design. For national governments, HOSIL offers a path to a sovereign cereal forecasting capability that is not hostage to any single external donor, which

Web of Agriculture: Journal of Agriculture and Biological Sciences
webofjournals.com/index.php/8

after 2025 is a matter of strategic prudence rather than convenience. For the humanitarian and development system, a credible, nationally owned forecast is exactly the input that anticipatory action and pre arranged finance need in order to function, and aligning the system outputs with the trigger, action, and financing structure of existing anticipatory action frameworks would let national forecasts plug directly into established disbursement mechanisms. For regional cooperation, a shared and transparent set of forecasts across neighboring countries that depend on the same crops and the same rivers can support coordinated response and can shift discussion from assertion toward measurement. And for the statistical system, the same pipeline strengthens official agricultural statistics by providing an independent, timely, spatially explicit verifying mechanism, advancing the long held goal of making agricultural statistics more sustainable and more granular.

Recommendations and a buildable roadmap

The following recommendations are prioritized, each justified, and together they form a phased roadmap for building HOSIL from the current state of the art.

First, prove the hybrid core on one crop in one region before broadening. Begin with winter wheat in a single set of provinces, establish that the data assimilated crop model plus a lightweight machine learning correction beats the climatological and vegetation index baselines at a several month lead, and only then expand. Breadth purchased before the core is proven is breadth wasted.

Second, treat calibration and trigger design as the central deliverable, not as a finishing touch. Because the entire value proposition rests on probabilities that mean what they say and on triggers that balance false alarms against missed events, invest disproportionately in calibration, in honest reporting of base rates, and in co designing triggers with the institution that will act on them.

Third, engineer sovereignty deliberately at every layer by choosing open and substitutable components, and budget for the national technical capacity to run them, because open inputs without local capacity is sovereignty in name only. The architecture should assume that any single external provider may disappear and should keep functioning when one does.

Fourth, connect the outputs to anticipatory action from the outset rather than as an afterthought. Express forecasts in the language of triggers, costed actions, and pre arranged finance, and engage the relevant disaster risk financing and anticipatory action partners early so that a fired trigger meets an existing disbursement mechanism rather than an empty channel.

Fifth, be explicit and disciplined about scope. State plainly that HOSIL forecasts the agricultural availability dimension of food security and that conflict, market, and access information must come from elsewhere, and design the food security interface to integrate those inputs rather than to substitute for them. Overclaiming completeness would be both scientifically wrong and ethically dangerous.

Sixth, keep a human in the loop and uncertainty on the surface. Treat every trigger as advisory evidence requiring authorization, present uncertainty as a first class output, and audit the equity consequences of action and inaction, so that the system protects the vulnerable rather than quietly transferring the cost of its errors onto them.

Seventh, plan for nonstationarity. Schedule continual recalibration, monitor for drift as the climate moves conditions outside the historical envelope, and lean on the mechanistic core for the unseen conditions in which warnings matter most.

Conclusion

The events of 2025 showed that the world capacity to see hunger coming was concentrated in a way that made it fragile, and that fragility is a choice rather than a necessity. The scientific ingredients for a different arrangement now exist: open satellite archives and open foundation models to observe the land, mature crop models and a hybrid methodology to forecast the harvest early and accurately even where labels are scarce, and an anticipatory action practice ready to convert a credible forecast into protection before the shock arrives. What has been missing is an integrated, sovereign, honest design that unites them. This paper has proposed one, named HOSIL after the harvest it exists to forecast, and has specified its architecture, its validation, and its failure modes without pretending it is finished. It is an architecture and a research agenda, offered in the conviction that a warning is only worth building if it is early enough to enable action, owned firmly enough to survive a donor shock, and honest enough about its own uncertainty that the people who act on it, and the people whose lives depend on that action, are not betrayed by false confidence.

References

1. Basso, B., Cammarano, D., Carfagna, E. (2013). Review of crop yield forecasting methods and early warning systems. Food and Agriculture Organization of the United Nations, Rome.
2. Becker-Reshef, I., Vermote, E., Lindeman, M., Justice, C. (2010). A generalized regression based model for forecasting winter wheat yields in Kansas and Ukraine using MODIS data. Remote Sensing of Environment, 114(6), 1312 to 1323.
3. Brown, C. F., Kazmierski, M. R., Pasquarella, V. J., et al. (2025). AlphaEarth Foundations: a geospatial embedding model and the Satellite Embedding dataset. Google and Google DeepMind. Earth Engine dataset GOOGLE/SATELLITE_EMBEDDING/V1/ANNUAL.
4. Defourny, P., Bontemps, S., Bellemans, N., et al. (2019). Near real time agriculture monitoring at national scale at parcel resolution: performance assessment of the Sen2 Agri automated system in various cropping systems around the world. Remote Sensing of Environment, 221, 551 to 568.
5. Famine Early Warning Systems Network (2025). Service interruption and resumption of operations. Reported by CNN, Devex, and the Center for Global Development, January to June 2025.
6. Jakubik, J., Roy, S., Phillips, C. E., et al. (2024). Prithvi EO 2.0: a versatile multitemporal foundation model for Earth observation applications. arXiv:2412.02732.
7. Kallenberg, M. G. J., Maestrini, B., van Bree, R., Ravensbergen, P., Pylianidis, C., van Evert, F., Athanasiadis, I. N. (2023). Integrating process based models and machine learning for crop yield prediction. arXiv:2307.13466.
8. Rojas, O. (2017). Early Warning Systems and crop yield estimation. In Handbook on Remote Sensing for Agricultural Statistics. Global Strategy to improve Agricultural and Rural Statistics (GSARS), chapter 6.
9. Wang, et al. (2024). Integrating data assimilation, crop model, and machine learning for winter wheat yield forecasting in the North China Plain. Agricultural and Forest Meteorology.
10. World Food Programme (2025). Disaster Risk Financing Annual Report; and Anticipatory Action programme materials. Rome.

11. Anticipation Hub (2024 to 2025). What is anticipatory action, and global anticipatory action statistics.
12. Global Agricultural Monitoring initiative (GEOGLAM) and the Agricultural Market Information System (AMIS). Crop Monitor for the major traded crops.
13. Food and Agriculture Organization of the United Nations, Global Information and Early Warning System (GIEWS) and the Agricultural Stress Index System (ASIS).

Web of Agriculture: Journal of Agriculture and Biological Sciences
webofjournals.com/index.php/8