EFRA

EXTREME FOOD RISK ANALYTICS

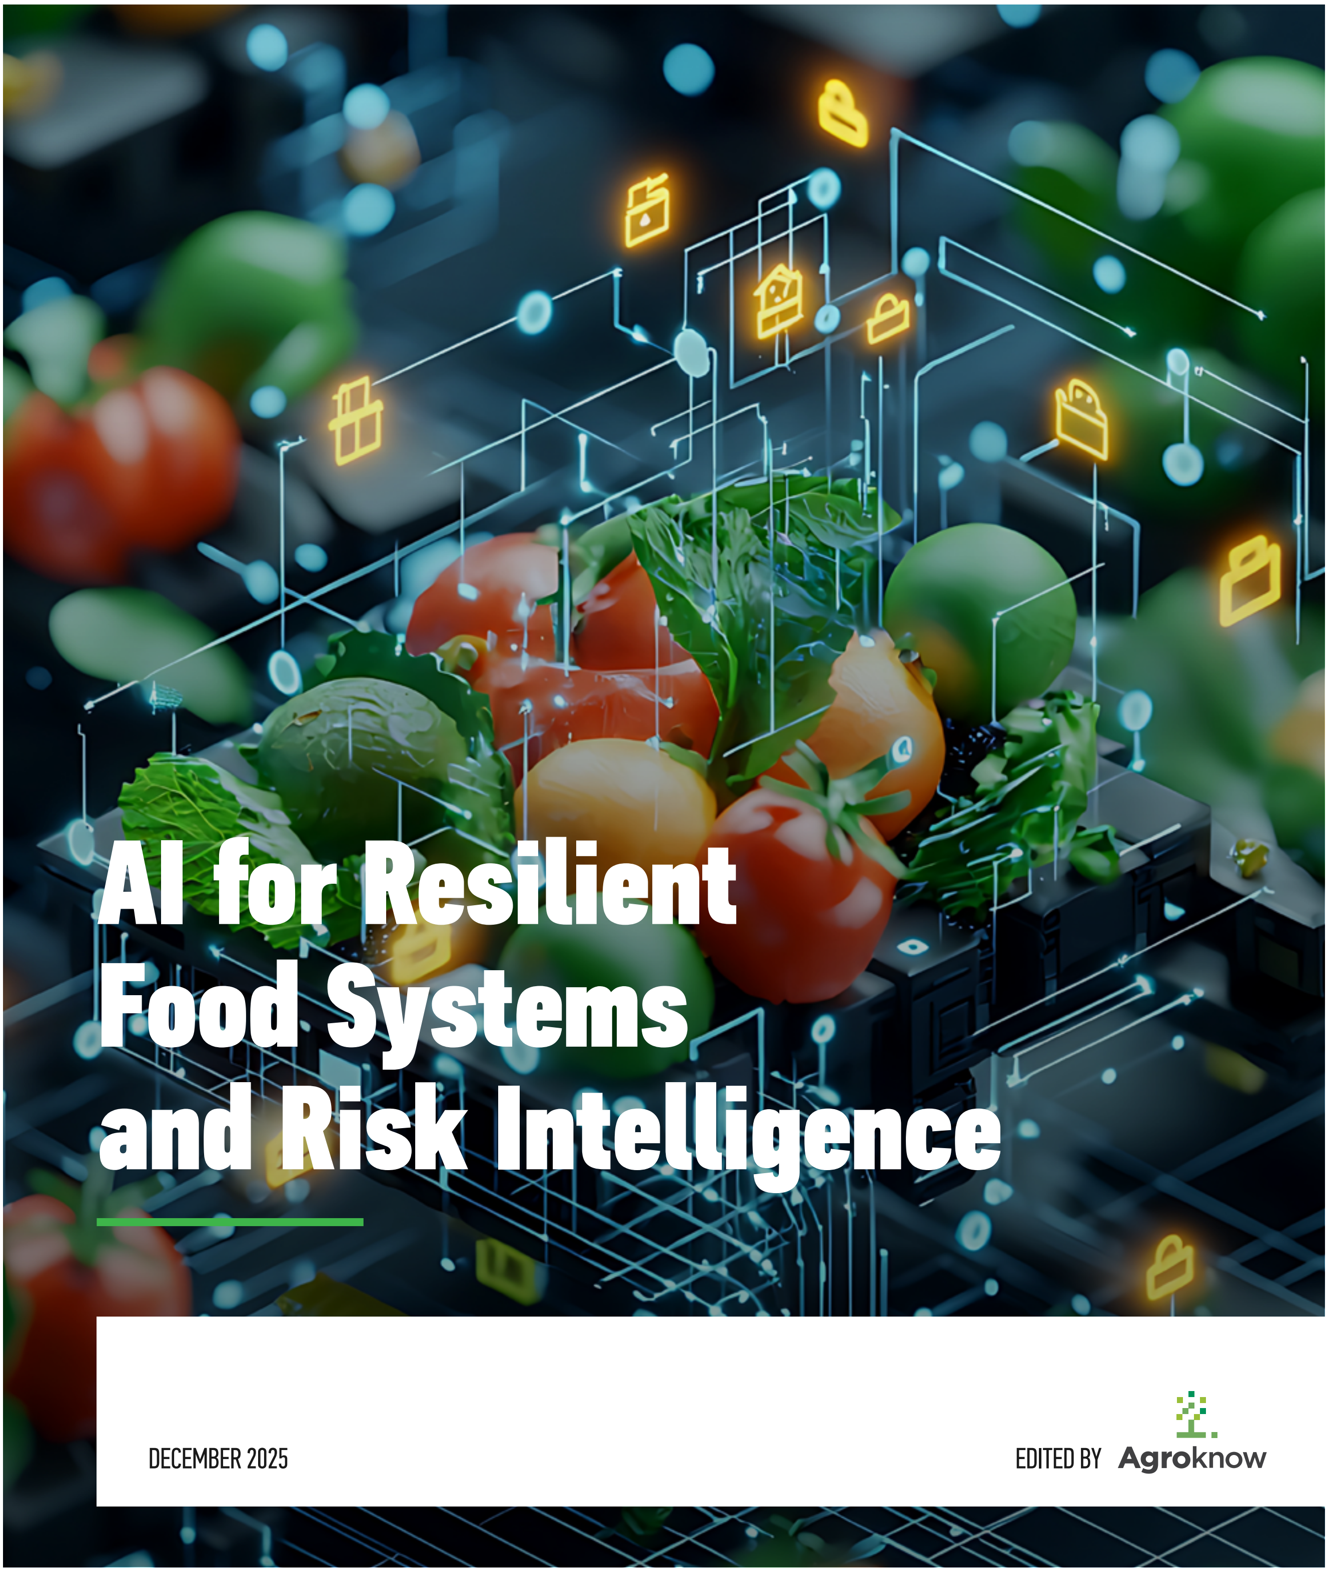

AI for Resilient Food Systems and Risk Intelligence

DECEMBER 2025

EDITED BY Agroknow

EFRA
EXTREME FOOD RISK ANALYTICS

Where trustworthy AI and food safety come together to strengthen Europe's resilience

This discussion paper was developed within the HORIZON EFRA project, which advances extreme data discovery, aggregation, and analytics to strengthen food-risk prevention and support resilient, trustworthy, and data-driven decision making across the food system. The project is coordinated by Dr. Babis Thanopoulos, Head of Innovation at Agroknow

TABLE OF CONTENTS

Foreword

Author
Prof. Saskia van Ruth
Professor of Food Supply Chain Integrity, University College Dublin, Ireland
Coordinator of the European Food Fraud Community of Practice (EFF-CoP)
Leader of the EU Cluster for Food Traceability and Trust

Let us begin from a food-system perspective. On a winter morning in March 2023, a batch of baby-leaf spinach left a greenhouse near The Hague, crossed two borders, and was expected to reach the Riviera within 30 hours. Forty-five minutes before unloading, an alert appeared on the logistics dashboard: a temperature excursion of 2.8 °C had occurred during transit through southern France, leading to accelerated bacterial proliferation and an estimated 18 % increase in the likelihood of Listeria growth. Within minutes, the shipment was diverted to a processor for blanching rather than entering the fresh-produce supply chain. The incident never reached consumers or the media. This near-miss illustrates a structural shift in food safety: artificial intelligence (AI) is increasingly enabling risks that once triggered recalls to be anticipated, detected, and mitigated in real time.

This discussion paper, AI for Resilient Food Systems and Risk Intelligence, brings together insights from Horizon Europe projects—including EFRA (Extreme Food Risk Analytics), HACID (Hybrid Human–Artificial Collective Intelligence in Open-Ended Decision Making), OASEES (Open Autonomous Programmable Cloud Applications and Smart Edge Sensors), and PLIADES (AI-Enabled Data Lifecycle Optimisation and Data Spaces Integration) alongside contributions from leading research organisations. Together, they examine how AI can strengthen food safety and system resilience through improved decision-making, early risk detection, cross-border coordination, and responsible deployment. They also articulate a forward-looking perspective on the role of AI in food systems towards 2035.

Across contributions, a common conclusion emerges: modern food systems must move beyond predominantly preventive and reactive models towards resilient systems capable of absorbing, adapting to, and acting upon disruptions before they escalate. The collection spans the full food chain, from primary production and precision agriculture to processing, logistics, climate services, and regulatory governance, demonstrating how AI can convert fragmented data into actionable risk intelligence. Contributors consistently emphasise predictive, explainable, and privacy-preserving AI as prerequisites for early warning, anticipatory intervention, and coordinated responses to transboundary threats. Recurrent themes include interoperable data spaces, digital twins and multimodal sensing for hazard forecasting, and the necessity of robust governance frameworks, human oversight, and shared standards to ensure trustworthy scaling.

Re-addressing the food system perspective, classic food-safety and authenticity management relies on management systems, retrospective epidemiology, periodic audits, and batch-based testing. While indispensable, these instruments are largely preventive or retrospective and offer limited capacity for real-time intervention. AI extends this framework in two complementary directions.

Firstly, early risk detection: AI models trained on multimodal sensor data—including spectroscopy, genomic outputs, and environmental variables—can detect deviations well before they become operationally visible. In the spinach example, a recurrent neural network could continuously evaluate temperature trajectories against pathogen growth models, updating the conditional probability of microbial outgrowth in real time.

Secondly, predictive analytics: AI systems integrating agronomic, climatic, trade, and behavioural signals can forecast disruptions such as mycotoxin outbreaks or logistical bottlenecks weeks in advance. Rather than identifying isolated failures, these models characterise disturbance pathways, allowing stakeholders to adjust sourcing, processing, or distribution strategies proactively.

A recurring insight in this collection is that resilience is not primarily a technological artefact, but an emergent property of connected systems. Yet food-system data remain highly fragmented: sensor data are proprietary, transactional data are siloed, and laboratory results are often locked in static formats. Several European projects address this challenge by making it possible to connect and use different types of data together, even when they come from incompatible systems.

The resulting transparency does more than improve analytics; it reshapes trust relationships among producers, regulators, and consumers. Resilience is realised when predictive insights trigger coordinated action across the supply chain, rather than remaining confined to individual actors.

At the same time, limitations need to be considered too. AI performance is constrained by data quality and representativeness; models optimised for European value chains may not transfer to smallholder or resource-limited contexts. Expanded sensor deployment raises concerns about e-waste, while large-scale computation carries energy costs. Most critically, asymmetric data ownership risks reinforcing existing power imbalances, particularly for actors at the beginning of the supply chain.

AI should therefore be understood not as an endpoint, but as a catalyst within a continuous cycle of sensing, anticipating, learning, and governance. The transition from reactive to resilient food systems is not merely a technical upgrade; it reflects a cognitive shift, i.e. from assumptions of stability to expectations of disruption, and from siloed optimisation to collective risk management. In this light, the spinach consignment that quietly changed course on a cold March morning represents more than an isolated logistics decision. It signals the emergence of a food system in which farms, vehicles, laboratories, and regulators are increasingly connected through shared intelligence, enabling risks to be addressed before they materialise. The contributions in this issue invite critical engagement with this transition: its methods, its limitations, and its ethical implications while advancing the shared goal of preventing tomorrow's food crises before they even begin.

Note from the Editor

Author
Dr. Babis Thanopoulos
Agroknow, Greece
EFRA Coordinator

Europe's food systems are entering a decisive decade. Climate volatility, globalised supply chains, emerging biological threats, and rapidly evolving production environments are reshaping how risks arise and how they must be managed. At the same time, the European Commission's Horizon Europe programme places strong emphasis on food safety, sustainability, data governance, and trustworthy Artificial Intelligence, calling for integrated, anticipatory, and science based approaches that strengthen resilience across the entire Farm to Fork continuum. This EFRA Discussion Paper brings together expert contributions from leading European universities, research centres, innovation projects, and food safety applications, offering a rich and multidisciplinary perspective on how Artificial Intelligence can support this transition.

These perspectives align closely with the mission of EFRA, which explores how extreme data mining, aggregation, and analytics can address the scientific, economic, and societal challenges associated with food safety and quality. EFRA's objectives, which include discovering and distilling food risk data from dispersed sources, designing human centred interfaces, demonstrating trustworthy, accurate, green, and fair Artificial Intelligence, and integrating big data, internet of things technologies, and advanced analytics, mirror the European Commission's priorities for data driven innovation, transparency, and responsible Artificial Intelligence adoption. EFRA's three pillars, the Data Hub, the Analytics Powerhouse, and the Data and Analytics Marketplace, reflect Horizon Europe's vision for interoperable data spaces, high impact analytics, and open innovation ecosystems where data holders, innovators, and decision makers collaborate to safeguard the food we eat.

A central message emerging from this Discussion Paper is the strong alignment between EFRA's work and the European Commission's priorities for Artificial Intelligence in food safety. The Commission emphasises the need for trustworthy, human centric, and transparent Artificial Intelligence systems, interoperable and sovereign European data spaces, early warning and predictive capabilities for emerging risks, and sustainable, resource efficient digital infrastructures that support the Farm to Fork and Green Deal objectives. EFRA's focus on extreme data discovery, explainable and privacy preserving analytics, cross border data interoperability, and human in the loop decision support directly responds to these priorities and contributes to the ambition of building a resilient, anticipatory, and science based food safety system for Europe.

Across the nine position statements, several shared themes emerge that resonate strongly with Horizon Europe priorities. Contributors emphasise the growing importance of predictive Artificial Intelligence for anticipating hazards before they escalate, whether through plant level disease detection, optical sensing, anomaly monitoring in robotics, climate risk intelligence, or hazard analysis and critical control point based risk forecasting. This shift from reactive to anticipatory risk management directly supports the Commission's goals for early warning systems, crisis preparedness, and climate adaptation. Many highlight the need for high quality, interoperable, and FAIR data, noting that fragmented datasets, inconsistent standards, and limited real time monitoring remain major barriers to effective risk prediction. Their insights align with the development of European level data infrastructures such as the Common European Agricultural Data Space, the Green Deal Data Space, and the European Open Science Cloud.

Trustworthiness is another recurring priority. Experts across domains stress that Artificial Intelligence systems must be transparent, interpretable, uncertainty aware, and auditable, especially when they support decisions with regulatory or safety implications. This reflects the requirements of the European Union Artificial Intelligence Act, which mandates traceability, human oversight, robustness, and fairness for high risk Artificial Intelligence systems. EFRA's focus on explainable, privacy preserving, and environmentally responsible Artificial Intelligence directly contributes to this European agenda. Several contributions also underline the importance of sustainable and energy efficient digital infrastructures, highlighting lightweight, resource efficient Artificial Intelligence architectures capable of operating reliably in constrained environments. This aligns with Horizon Europe's commitment to green digital transformation and the need to ensure that Artificial Intelligence adoption does not increase environmental burdens.

Cross border cooperation and federated intelligence emerge as essential components of future food system resilience. Food safety risks and climate hazards do not respect national boundaries, and contributors highlight the need for privacy preserving analytics, federated learning, and interoperable governance frameworks that enable collaboration without compromising data sovereignty. These ideas reflect the Commission's priorities for European data spaces, cross border early warning networks, and coordinated risk management. Several statements also emphasise the importance of hybrid human and Artificial Intelligence supported decision making, combining computational power with human expertise, in line with the European Union's vision for human centric Artificial Intelligence.

Looking ahead to 2035, the visions presented in this paper are ambitious yet grounded. They include real time sensor networks for precision agriculture, hybrid human and Artificial Intelligence systems for climate risk intelligence, federated European food data spaces, multimodal digital twins of the food system, and cross sector Artificial Intelligence ecosystems capable of detecting anomalies, forecasting disruptions, and coordinating responses across borders. These visions reflect a shared belief that resilience will increasingly depend on anticipatory, data driven, and collaborative intelligence supported by trustworthy Artificial Intelligence and robust governance frameworks.

Together, the nine contributions offer a compelling roadmap for how Europe can harness extreme data analytics and Artificial Intelligence to strengthen food system resilience. They demonstrate that progress will depend not only on technological innovation, but also on governance, interoperability, sustainability, and human centred design, all of which are central pillars of Horizon Europe. As EFRA continues to engage public and private stakeholders, the collaborative spirit reflected in this Discussion Paper will be essential for realising the full potential of Artificial Intelligence in safeguarding Europe's food systems. By aligning scientific excellence with European values of trust, transparency, sustainability, and fairness, EFRA contributes to a future where food safety is proactive, data driven, and resilient by design.

AI METHODS &
AUTONOMOUS SYSTEMS

AI METHODS & AUTONOMOUS SYSTEMS

National Technical University of Athens (NTUA)

Authors:
Professor **Konstantina S. Nikita**
Maria Athanasiou, Postdoctoral Researcher
Biomedical Simulations and Imaging Laboratory,
School of Electrical and Computer Engineering (BioSim)

Where can AI add the most value in food safety decision-making?

AI offers the greatest value where decisions depend on integrating diverse, rapidly evolving information streams. In the context of food safety, decision-making is hindered by multi-layered global supply chains, heterogeneous data formats, dynamic environmental and market conditions, regulatory variation, and multifactorial risks that interact in complex ways. Weak or diffuse early signals often go unnoticed because surveillance still relies on manual reporting, isolated laboratory results, or static risk assessments unable to reflect real-time system variability. As a result, critical operational and regulatory tasks, including prioritising inspections, tracing contamination pathways, validating supplier compliance, assessing supply-chain vulnerabilities, or triaging high-risk products, remain slow and reactive. Data fragmentation, inconsistent standards, limited real-time monitoring, sparse contextual metadata, and poor integration across environmental, microbiological, trade, and behavioural datasets further reduce situational awareness and restrict the ability to anticipate cascading disruptions across the food system.

AI can bridge these gaps by fusing multimodal data into coherent early-warning intelligence. Techniques including graph-based models can map and propagate risk across supply chains, deep learning can detect anomalies or foreign objects, and probabilistic, uncertainty-aware models can indicate when interventions are needed. Generative and representation-learning approaches can enrich sparse datasets and reveal latent risk factors, while the combination of mechanistic knowledge with machine learning methods can guide the development of risk-informed strategies, providing scenario analysis for climate-driven hazards, geopolitical disruptions, emerging pathogens, or supply-chain stressors. At the operational level, AI-driven pipelines support

Position Statement from BioSim

automated food classification, quality grading, fraud detection, and contamination screening, reducing dependence on resource-intensive laboratory testing. By enabling earlier interventions, reducing economic losses from outbreaks and recalls, and improving overall system resilience, AI transforms food safety from reactive incident management into anticipatory, risk-intelligent decision-making, grounded in timely, evidence-based insights that benefits regulators, producers, and consumers alike.

How is AI helping detect and respond to emerging risks in your sector?

Across the health sector, AI has become a central tool for strengthening early detection, surveillance, and adaptive response. Its role has expanded from traditional analytics to sophisticated applications, including real-time outbreak monitoring, risk prediction, personalised diagnostics, and treatment optimisation. The recent acceleration of digital innovation, driven in part by the COVID-19 pandemic, has brought forward cutting-edge approaches including generative AI agents for rapid evidence synthesis, knowledge graphs for linking heterogeneous biomedical data, continual learning systems that adapt to evolving conditions, and increasingly mature explainable and trustworthy AI methodologies.

Within this broader landscape, our research focuses on disease modelling and developing reliable and robust adaptive frameworks for the prognosis, diagnosis, and management of complex diseases. In this direction, we fuse multimodal clinical, physiological, behavioural, molecular, and environmental data using physiology-informed mathematical models, advanced machine learning and deep learning, interpretability techniques, uncertainty quantification, and bias mitigation and domain adaptation methods. These approaches address critical challenges related to the clinical adoption of AI-based systems - trustworthiness, interpretability, generalisability, and fairness - and have been applied to diabetes, cardiovascular conditions, neurodegenerative diseases, cancer, and COVID-19, demonstrating strong capability in detecting evolving and multimodal risk signals.

For example, our LSTM-based influenza-like illness forecasting system combined surveillance data, weather conditions, and Twitter activity and showed that multimodal fusion significantly outperforms single-source models. Similarly, our drift-adaptive COVID-19 detection framework, our interpretable Alzheimer’s disease diagnostic models, and our uncertainty-aware adaptive frameworks for cardiovascular risk stratification illustrate how integrating heterogeneous data with advanced AI-driven modelling approaches enhances detection sensitivity and adaptive response. These same principles are directly transferable to food safety for identifying emerging contamination routes, environmental stressors, supply-chain anomalies, or behavioural drivers of risk in real time.

What are the main challenges in deploying AI for resilient food systems?

AI deployment for resilient food systems is constrained by a combination of technical, organisational, and governance challenges. On the technical side, data remain highly fragmented and heterogeneous, originating from microbiological analyses, sensor networks, climate models, logistics records, and behavioural sources, often with incompatible formats, inconsistent ontologies, and limited contextual metadata. This lack of interoperability makes it difficult to integrate information across the supply chain, while sparse real-time monitoring and uneven data quality hinder the development of reliable and adaptive early-warning systems. Models often struggle with generalisability as food systems evolve - changes in supply-chain behaviour, pathogen dynamics, or

Position Statement from BioSim

environmental conditions can trigger performance drift. In this context, the lack of mechanisms for drift detection and model adaptation results in degraded performance as patterns shift over time and/or across domains.

Trustworthiness is another key barrier: although regulators and industry stakeholders require interpretability and confidence estimates to assess the reliability of model outputs, many current tools remain opaque, providing point estimates without transparent reasoning or confidence measures. At the organisational level, uneven digital maturity, limited analytical expertise, and the difficulty of embedding AI into existing inspection and quality-assurance workflows hinder effective adoption. Governance challenges further pronounce these issues, particularly around regulatory compliance, as organisations must navigate evolving requirements, including robustness, fairness, accountability, and human oversight under the EU AI Act, while balancing data-sharing constraints linked to confidentiality, commercial sensitivity, and GDPR.

Progress in resilient food systems depends on domain-specific infrastructures, trustworthy modelling, and collaborative governance. Harmonised standards and practical data-sharing mechanisms must enable safe exchange across the supply chain, while drift detection, interpretability, and fairness are built into AI design.

What does trustworthy AI look like in the context of food safety?

Trustworthy AI is built on four foundational pillars - interpretability, uncertainty awareness, robustness, and fairness - supported by regulatory compliance and meaningful human oversight. These principles ensure that AI systems behave transparently, reliably, and safely in high-stakes environments. In food safety, interpretability means that models must provide clear, traceable explanations of risk alerts so inspectors and regulators can understand the underlying evidence. Interpretability techniques widely adopted in various domains including feature attribution, surrogate models, counterfactual explanations, and rule-based components, can be directly applied to make food-safety predictions understandable and auditable.

Uncertainty awareness ensures that AI systems communicate how confident they are in a prediction, allowing authorities to distinguish high-confidence alerts from cases where human review is needed: Bayesian inference, ensembles, or Monte Carlo dropout provide explicit confidence measures that prevent over- or under-reaction. Robustness requires models to remain reliable as supply-chain patterns, pathogens, or environmental drivers evolve. Various approaches for drift detection, continual learning, and domain adaptation have been proposed to address such challenges and can be leveraged in the field of food safety.

Bias-mitigation strategies, such as balanced sampling, bias audits, algorithmic modifications, and counterfactual analyses, can help prevent systematic disparities in food-safety decisions, ensuring consistent performance across regions, production systems, and product categories. By operationalising these pillars through well-validated methodologies, AI systems in food safety can become transparent, adaptive, scientifically grounded, and suitable for regulatory and industry use.

How can AI support cross-border coordination and data sharing to enhance food system resilience?

AI can strengthen cross-border coordination by enabling data harmonisation, shared digital infrastructures, secure information exchange,

Position Statement from BioSim

shared early-warning intelligence, and coordinated response mechanisms. AI systems can harmonise heterogeneous datasets through automated ontology mapping, knowledge graphs, and multimodal fusion, allowing environmental, microbiological, trade, and logistics data from different countries to become interoperable. Progress depends on common data standards and shared infrastructures, including unified terminologies, metadata schemas, reporting formats, and connectivity frameworks that allow AI tools to operate reliably across jurisdictions.

Effective coordination also requires digital readiness, ensuring all countries have the capacity, governance structures, and technological foundations to engage in cross-border data sharing. Privacy-preserving technologies such as federated learning and differential privacy can enable collaborative model development without exposing sensitive commercial or personal data. AI-driven traceability systems, powered by IoT sensors, blockchain, and digital product passports, are able to enhance real-time cross-border visibility and support rapid, coordinated recalls.

Shared early-warning models that integrate climate indicators, pathogen dynamics, supply-chain signals, and trade flows can detect transnational threats and support joint assessment and response. Underpinning these capabilities are collaborative governance frameworks that define responsibilities, ensure accountability, support safe data sharing, and institutionalise long-term cooperation. Together, these components enable an integrated, intelligence-driven approach to food-system resilience at regional and global scale.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems?

By 2035, a transformative breakthrough would be the adoption of adaptive multimodal risk-intelligence platforms that can reliably forecast emerging food-system threats before they materialise. Such platforms would integrate sensor data, laboratory findings, climate signals, and supply-chain dynamics into a unified digital environment, powered by explainable models that continuously recalibrate as conditions evolve. Rather than static dashboards, this would resemble a dynamic digital twin of the food system, capable of running scenario simulations, stress-testing vulnerabilities, and guiding targeted interventions with traceable reasoning.

Realising this vision will depend not on producing more AI tools, but on building the foundations that make such platforms reliable: high-quality and interoperable data, robust, transparent model-validation frameworks, and regulatory capacity to evaluate adaptive, uncertainty-aware systems. Experience from healthcare - where multimodal fusion, advanced AI-driven analytics, drift-adaptive diagnostics, and rigorous validation already support complex decision-making - shows how such a paradigm can be made operational. Bringing this approach to food safety would enable a shift from episodic surveillance to predictive resilience, empowering decision-makers with anticipatory, scientifically grounded intelligence.

AI METHODS & AUTONOMOUS SYSTEMS

HORIZON HACID Project

National Research Council (CNR)

Author:
Vito Trianni
Institute of Cognitive Systems and Technologies (ISTC-CNR)
HACID Coordinator

Met Office

Authors:
Anrijs Abele
Neha Mittal
Fai Fung

Where can AI add the most value in food safety decision-making?

Climate change is exerting profound pressure on global food systems, exposing them to increasingly frequent and severe hazards such as prolonged droughts, extreme rainfall, soil degradation, and the emergence of new pests and pathogens. These disruptions reveal a structural vulnerability: traditionally, food systems have relied on reactive approaches, responding to shocks only after they materialize.

In recent years, climate services have emerged as a domain dedicated to translating climate science into actionable information and tools that support informed, anticipatory decision-making. At the same time, artificial intelligence (AI) is reshaping the capacity of climate services to move from narrow, physical science-based analysis to comprehensive climate risk intelligence. The HACID project (https://www.hacid-project.eu) offers a timely example of how hybrid human–AI systems can support more resilient decision-making in this evolving landscape.

Climate services can be broadly understood as the processes and tools that transform raw climate data into actionable, context-specific information for decision-makers. This includes the production of datasets based on observations and simulated futures at seasonal to multi-decadal timescales which are then interpreted, translated and communicated to a wide range of audiences.

For food systems, climate services serve an increasingly critical role by supporting farmers, policymakers, and supply-chain actors in managing both immediate and long-term risks. Seasonal forecasts currently inform planting choices and irrigation planning; climate projections guide the design of resilient infrastructure, cultivation of climate resilient crop varieties and insurance schemes; and hazard analyses illuminate how climatic shifts may affect crop suitability, market stability, and food safety. In essence, climate services offer a bridge between complex climate science and the diverse operational and long-term strategic decisions required to safeguard food system functioning.

How is AI helping detect and respond to emerging risks in your sector?

The introduction of AI into climate services strengthens this bridge in several important ways. First, AI enhances the capacity to synthesise the immense and rapidly expanding volume of climate data produced by observational systems and climate models. The latest CMIP6 ensemble alone generates datasets so large and complex that making sense of the available information is increasingly impractical.

Second, AI supports predictive modelling across multiple timescales, as well as downscaling global and regional projections to obtain high-resolution information. More accurate data means improving early detection of risks such as drought onset, detecting patterns like shifts in precipitation regimes, or detecting relevant conditions conducive to crop diseases. These capabilities allow climate information to be used not only descriptively but also prognostically, enabling earlier and more targeted interventions.

Third, generative AI contributes to decision support by structuring complex adaptation workflows, exploring alternative scenarios, and helping users evaluate the robustness of adaptation strategies under uncertainty. When combined with human expertise, AI can help reduce cognitive burdens, limit bias, and increase the completeness of assessments.

What are the main challenges in deploying AI for resilient food systems?

Despite these contributions, the deployment of AI in climate services is not without challenges. The fragmentation of climate, agricultural, and socioeconomic datasets - both in terms of availability and compatibility - continues to impede integrated analysis, while issues of transparency and explainability create barriers to trust, especially when AI-generated outputs influence high-stakes decisions.

The diversity of potential users - from smallholder farmers to national/international policy makers - means that climate services must be tailored to differing capacities and contexts, raising questions about accessibility and usability. Ethical considerations, including bias embedded in data and the risk of privileging the needs of data-rich regions, further complicate deployment.

Moreover, because food systems cross national borders, the absence of harmonised governance frameworks limits the ability of AI-driven climate services to operate effectively at the transnational scale required by global supply chains.

What does trustworthy AI look like in the context of food safety?

The HACID project (Hybrid Human-Artificial Collective Intelligence for Decision Support in Open-Ended Domains) directly addresses several of these challenges by exploring how collective intelligence - emerging from collaboration between human experts and AI agents - can support climate-related decisions. HACID focuses specifically on climate services, providing a decision support system (DSS) that helps policymakers and organizations adapt to uncertain future climate conditions. Central to the project is the construction of an extensive domain knowledge graph (DKG) that integrates climate projections, datasets endorsed by national governments, and established methodologies for selecting and processing climate information. This DKG maps the relationships among climate models, hazards, indices, and methodologies for climate information analysis, thus formalizing complex reasoning processes that climate experts typically undertake.

Within the HACID DSS, experts confronted with a specific problem - such as change in surface water flooding risk - identify and annotate relevant elements of the DKG, drawing on their disciplinary knowledge, and propose workflows that can lead to the best risk assessment. AI agents operate in parallel, proposing their own structured solutions based on the knowledge resources, potentially exploiting the wide diversity of methods and approaches modelled in the DKG. The system then synthesizes insights from all contributors, producing a more comprehensive and transparent pathway for extracting relevant climate information. This hybrid approach enhances early detection of emerging risks by surfacing connections that might be overlooked by individual experts and increasing the diversity of approaches, while also providing an auditable rationale for decision-making. It thus supports a more resilient form of climate service provision, grounded in both computational power and human interpretative capacity.

How can AI support cross-border coordination and data sharing to enhance food system resilience?

Extending HACID to food safety and food system resilience would require several next steps. The knowledge base would need to incorporate climate-sensitive food safety risks, such as waterborne pathogen dynamics, contamination pathways, and the effects of extreme weather on storage environments, as well as food system resilience, such as characterisation of abiotic and biotic factors that influence crop productivity.

Strengthening cross-border coordination would require interoperable governance frameworks to ensure the secure, trustworthy exchange of climate and food safety data. Such integration is essential because climate hazards that affect food safety do not respect national boundaries, and supply chains increasingly depend on synchronised risk communication.

Moreover, embedding food safety regulatory processes into the HACID DSS would help ensure that recommended adaptation strategies align with existing inspection protocols, legal frameworks, and operational workflows across different jurisdictions. This alignment would enhance trust in AI-supported decisions and facilitate coordinated action across the food system.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems?

Looking ahead to 2035, a transformative breakthrough would be the emergence of a globally integrated, AI-enabled climate-risk intelligence infrastructure dedicated to food systems. Such an infrastructure would combine long-term climate projections with real-time monitoring from sensors, satellites, and supply-chain data streams to provide risk narratives.

It would be capable of continuously detecting emergent risks and opportunities while simulating alternative adaptation pathways and recommending context-appropriate actions. Importantly, this system would operate through hybrid human–AI teams, ensuring that scientific rigour, local knowledge, and ethical considerations remain central.

Achieving such an infrastructure would represent a decisive shift from reactive crisis management to anticipatory governance of climate-resilient food systems, aligning closely with the vision that HACID is beginning to articulate.

AI METHODS & AUTONOMOUS SYSTEMS

HORIZON PLIADES Project

Centre for Research and Technology Hellas (CERTH)

CERTH
CENTRE FOR RESEARCH & TECHNOLOGY HELLAS

Authors:
Dimitrios Giakoumis, Researcher
Kosmas Tsiakas, Research Assistant
Theodora Kontodina, Research Assistant
Ioannis Mariolis, Research Assistant
Information Technologies Institute (ITI)
PLIADES Coordinator

Where can AI add the most value in food safety decision-making?

Although PLIADES does not target the food domain directly, its architecture for full-data-lifecycle optimisation and AI-enabled interoperable data spaces demonstrates how complex distributed data ecosystems can support evidence-based decision-making. In food safety contexts, similar approaches could be adapted to connect inspection records, IoT sensor data, and logistics tracking under shared semantics and governed access rules. By enabling real-time linkage of heterogeneous datasets, the PLIADES framework offers a transferable model for shifting from reactive to proactive food safety management, by reducing the latency between data acquisition, analysis, and strategic or operational decisions.

How is AI helping detect and respond to emerging risks in your sector?

PLIADES develops AI-assisted mechanisms for data quality monitoring, semantic alignment, and integrity validation within interconnected and interoperable data spaces. These capabilities ensure that information exchanged across sectors remains consistent, reliable, and actionable - conditions essential for early risk detection and response. While PLIADES applies these methods in mobility, healthcare, manufacturing, energy, robotics, and Green Deal use cases, the same principles could be extended to food systems, where identifying risks such as contamination, fraud, or supply-chain disruptions depends on timely, high-quality data. By enabling federated analytics and trusted cross-domain data sharing, PLIADES provides a transferable framework for proactive, AI-driven risk intelligence.

What are the main challenges in deploying AI for resilient food systems?

The main barriers in deploying AI for resilience in food systems (many of which are tackled in

Position Statement from CERTH / PLIADES

PLIADES) span technical, organisational, and regulatory dimensions. Across its mobility, healthcare, manufacturing, energy, robotics, and Green Deal domains, PLIADES addresses data fragmentation, lack of interoperability standards, and trust deficits that limit cross-sector data use. It develops semantic alignment methods, data quality frameworks, explainable AI tools, and governance models to ensure transparency and sovereignty. These enablers, together with shared infrastructures and capacity building for smaller actors, are equally relevant to food systems, where reliable, explainable, and ethically governed AI is essential for safe and adaptive decision-making.

What does trustworthy AI look like in the context of food safety?

In PLIADES, trustworthiness is achieved through transparency, explainability, and continuous human oversight across the AI and data lifecycle. The project develops metadata-driven traceability, provenance tracking, and explainable AI methods that make automated insights auditable and interpretable by end users. Every AI-assisted decision can be linked to its data origin, validation status, and confidence level, supporting accountability and human-in-the-loop supervision. While applied in mobility, healthcare, manufacturing, energy, robotics, and Green Deal domains, these principles can inspire trustworthy AI in food safety - where clear reasoning, transparent data flows, and ethical oversight are essential for building confidence among regulators, producers, and consumers.

How can AI support cross-border coordination and data sharing to enhance food system resilience?

PLIADES advances a federated architecture of interoperable European data spaces, where AI enables collaboration without compromising data sovereignty or privacy. Through semantic harmonisation, metadata registries, and trusted governance frameworks, the project demonstrates how AI can align standards and support shared intelligence across borders and sectors. Applied to food systems, such an approach could enable coordinated responses to emerging risks by linking regional data under common semantics and ethical rules. PLIADES' work across mobility, healthcare, manufacturing, energy, robotics, and Green Deal domains provides a reusable blueprint for resilient, cross-border data ecosystems compliant with EU data-space and AI policy priorities.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems?

By 2035, the vision inspired by PLIADES is a federated European Food Data Space where AI continuously monitors, analyses and predicts risks across the entire food-value chain. Data from farms, processing plants, logistics, and regulators would interconnect through trusted, interoperable infrastructures guided by transparent governance models. Such a system would enable real-time, cross-border collaboration, where AI not only detects emerging threats but anticipates them through continuous learning and multi-sector data integration. Building on the interoperability, governance, and trust mechanisms developed in PLIADES, this future ecosystem would make resilience an inherent property of Europe's food systems.

AI METHODS & AUTONOMOUS SYSTEMS

HORIZON RoboSAPIENS Project

Aristotle University of Thessaloniki (AUTH)

Authors:
Professor **Anastasios Tefas**
Vasilios Moustakidis, Ph.D Student
Department of Informatics
Assistant Professor **Nikolaos Passalis**
Department of Chemical Engineering
RoboSAPIENS Project Partner

Aarhus University

Author:
Professor **Peter Gorm Larsen**
Department of Electrical and Computer Engineering
RoboSAPIENS Project Coordinator

How can predictive AI technologies strengthen resilience in food systems or related domains?

Predictive AI becomes far more reliable when supported by a structured understanding of where anomalies can emerge within an AI-powered system. These anomalies can arise not only from sensors or hardware, but also from the behaviour of the AI model itself, for example, when it encounters situations, it was not trained on or when its internal representations drift over time. By examining the full decision pipeline, from sensing to execution, organisations can recognise early signs that the system is diverging from expected behaviour. Rather than assuming complete knowledge of how anomalies emerge, modern approaches focus on detecting deviations from normal operational patterns, even when those deviations have never been observed before. This strengthens resilience in domains such as robotics and can be transferred to food systems, where early identification of unexpected readings, equipment irregularities, or process fluctuations is essential for safety and continuity. Combining predictive models with continuous monitoring of system behaviour helps organisations anticipate risks sooner, reduce uncertainty, and maintain stable and trustworthy operations.

Position Statement from RoboSAPIENS

What technical or organisational barriers limit the deployment of predictive AI in food systems or similar contexts?

Several factors limit the effective deployment of predictive AI and anomaly detection in critical systems. A key technical challenge is that AI models must not only interpret their environment but also estimate the uncertainty within it. Building models capable of recognising both expected signals and uncertain or ambiguous situations requires development methodologies that are more complex than those used in traditional prediction tasks. At the system level, the surrounding framework must also be able to handle uncertainty: the decision-making pipeline should integrate not only the model's outputs but also information about confidence levels or potential anomalies arising from different components. When we deal with robots we need to make sure that they remain safe for the objects (including humans) in their environment, even when parts of their control have been trained using AI components. This demands continuous monitoring, real-time processing, and infrastructures able to combine signals from sensing, interpretation, and planning stages. Interoperability issues, such as incompatible data formats or isolated software modules, further limit the ability to form a unified view of system behaviour. Organisational factors also contribute: limited familiarity with AI-based diagnostics, uncertainty about accountability when automated alerts are raised, and reluctance to modify established procedures can slow adoption. Together, these barriers make it challenging to build dependable, transparent predictive systems capable of detecting and responding to anomalous behaviour early.

What are the key enablers and barriers for data sharing in the food industry, and how can AI help?

Sharing operational and monitoring data is important for detecting anomalies early, as it helps create a clearer picture of how a system behaves across different stages of sensing, interpretation, and decision-making. Ensuring the safety of the decisions taken is paramount for the trustworthiness of any system. However, several factors limit this kind of sharing. Organisations may be cautious about sharing internal system signals due to privacy concerns, unclear responsibilities, or uncertainty about how diagnostic information will be used. Technical barriers also play a role: incompatible formats, isolated tools, and fragmented infrastructures prevent the integration of data needed to understand unusual behaviour in context. AI can support more secure collaboration through privacy-preserving methods, shared representations that protect sensitive details, and automated checks that ensure data quality. In addition, AI-on-the-edge and embodied intelligence can reduce the need to share raw data externally by enabling systems to analyse signals locally and transmit only essential insights. These capabilities help create reliable foundations for early anomaly detection without compromising confidentiality or operational security.

How do you approach explainability and trust in AI systems used for risk prediction?

Building explainability and trust in AI systems begins with making their internal behaviour visible and interpretable. A practical approach is to examine how signals evolve across the different stages of an AI system, such as perception, interpretation, and decision-making, to identify when the system starts to diverge from expected behaviour. Highlighting these changes in clear, human-readable formats helps operators understand why a particular decision was made and when an intervention may be necessary. Trust is further supported through mechanisms that allow experts to review alerts, validate the system's reasoning, and adjust parameters when unusual behaviour is detected. By combining transparent monitoring with human oversight, AI systems used for risk prediction can provide more dependable and understandable outputs, even in complex operational environments.

What considerations guide the design of sustainable AI infrastructure in your organisation or field?

Sustainable AI infrastructure prioritises efficiency, stability, and long-term reliability. In robotics and autonomous systems, this means designing models that can monitor behaviour and detect anomalies without requiring excessive computational resources to guarantee safety. Lightweight architecture reduces energy consumption and makes continuous operation more feasible, especially in environments where systems must run safely and reliably for extended periods. Sustainability also involves creating models that remain stable over time, limiting the need for frequent retraining or manual adjustments. It is equally important to develop interfaces and model architectures that can be upgraded directly, without extensive hardware changes or major redesigns of the decision-making pipeline. Standardised frameworks help support this by enabling components to be updated, replaced, or extended with minimal disruption. These considerations are relevant to food systems as well, where energy-efficient, easily maintainable, and dependable AI tools support resilient operations and reduce the environmental footprint of data-driven processes.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems or cross-sector collaboration?

By 2035, a breakthrough would be the widespread adoption of AI systems that not only monitor their own behaviour continuously but also adapt effectively to changes in their environment. Beyond identifying irregularities future systems should be able to update their internal models, adjust parameters, and refine their responses as conditions evolve, ideally without requiring complete retraining or disruptive recertification processes in critical sectors. Building on advances in anomaly detection and system understanding, next-generation AI should be able to:

- detect unusual behaviour early and respond safely,
- adjust to changing environments while maintaining dependable performance,
- provide clear explanations when something goes wrong, and
- follow established standards for trustworthy and transparent operation.

Such capabilities would allow robotic platforms, food-processing equipment, and supply-chain technologies to operate more reliably, even under uncertainty. Achieving this vision will require strong collaboration across research, industry, and policy. Enhancing early anomaly detection and coordinated response mechanisms is essential for creating AI systems that society can depend on in critical domains where safety for their environment also needs to be guaranteed.

FOOD SAFETY & AGRICULTURE APPLICATIONS

FOOD SAFETY & AGRICULTURE APPLICATIONS

HORIZON OASEES Project

National Centre for Scientific Research "DEMOKRITOS" (NCSR Demokritos)

DEMOKRITOS

Author:
Dr. Akis Kourtis, Researcher
Institute of Information & Telecommunications (IIT)
OASEES Project Coordinator

Where can AI add the most value in food safety decision-making?

Food safety decision making is a critical sector for European citizens, where complexity and diverse supply chains create a challenging task for reliable tracking of food sources and a trustworthy monitoring system for the different stakeholders involved. AI as an enabler via the integration of different trust based technologies, i.e., blockchain, can not only accelerate automation of different processes, but also increase trust and traceability of the entire action chain.

In this respect, OASEES extends the blockchain based paradigm of Decentralized Autonomous Organizations (DAOs), where decision making is governed by smart contracts, which are transparent and controlled by its members, which can be both humans and AI agents. This convergence creates an interface for experts/regulators to approve or override different AI based decisions, on zero-trust policy. AI adoption may face obstacles based on trust issues, therefore a systematic approach on traceability and auditing of decisions can truly benefit such a critical sector as is food safety of today.

How is AI helping detect and respond to emerging risks in your sector?

Edge inference and data processing close to the source can be a significantly beneficial factor in different aspects of the food supply chain, as they not only accelerate procedures but also preserve the privacy and integrity of the data produced and inferred.

In this respect, OASEES follows an edge-first approach in its programmable framework, adopting certain primitives of the compute-to-data paradigm. The AI models produced are deployed close to the site, with edge processing tailored for

Position Statement from IIT / OASEES

resource-constrained devices (e.g. swarms), which can process data ad hoc, limiting cloud resource usage to the bare minimum.

From the perspective of food systems, the proposed paradigms can be directly applied, since the amount of generated data is vast, and it is logical to constrain data processing close to the source.

What are the main challenges in deploying AI for resilient food systems?

A key challenge for AI deployment in food systems is heterogeneity. The number of different cases and their corresponding enablers varies significantly, creating diverse requirements and limitations, which makes it difficult to establish a holistic approach.

This also creates challenges from a scaling perspective, since scalability normally depends on unified resource management, especially within the cloud/edge continuum. Different cases scale differently, and a horizontal approach - particularly in a resilient food system - can generate a multitude of challenges in building a fully functional end-to-end lifecycle.

What does trustworthy AI look like in the context of food safety?

For food safety, this translates into trusted AI pipelines where data provenance is verifiable across the continuum, risk scores and recommendations are traceable to models and data sources, and regulators or quality managers can inspect DAO records showing how alerts were handled, which thresholds were changed, and who authorised each step.

OASEES proposes a layered architecture with segregation of security zones and data-minimising designs, supporting compliance-friendly, "explainable-by-design" AI. Explainability and human oversight are realised via human-in-the-loop mechanisms embedded in DAO workflows, where domain experts validate data, vote on actions, and review logs of robot or service behaviour.

How can AI support cross-border coordination and data sharing to enhance food system resilience?

OASEES aligns with European initiatives for sovereign, interoperable data spaces (e.g. Gaia-X, IDSA) and demonstrates how federated operation across multiple operators and jurisdictions can be implemented.

In the food sector, this could underpin cross-border early-warning networks where AI models run locally on national infrastructures but share anonymised features, risk indicators, or aggregated traces through governed data spaces - aligning standards while respecting local rules and commercial sensitivities.

Furthermore, it enables participants to share AI-ready data products under explicit policies while retaining sovereignty.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems?

A desirable 2035 breakthrough is a pan-European "food-risk swarm continuum": thousands of interoperable swarms of sensors, robots, logistics nodes, and analytical services, each operated by different actors but orchestrated via common SDKs, DAOs, and data-space rules.

In such a system, edge AI would continuously assess contamination, fraud, and infrastructure risks; service-mesh technologies would automatically reconfigure flows and capacity; and DAO-based governance would align incentives so that farmers, processors, retailers, and authorities collaboratively manage risk.

Rather than isolated pilots, this would be a living, self-optimising network where every new sensor or model instantly strengthens collective foresight, and where transparency, accountability, and human oversight are embedded by design.

FOOD SAFETY & AGRICULTURE APPLICATIONS

Digital Quality & Food Safety Consultancy (DQFS Consultancy)

Author:
Geert van Kempen, Owner & Principal Advisor

How can predictive AI technologies strengthen resilience in food systems or related domains??

Predictive AI can strengthen food system resilience by transforming digitally structured food safety data into probabilistic risk intelligence. When hazard analyses, process controls, and supply parameters are encoded in interoperable formats, such as those emerging from digital HACCP systems, machine learning models can detect precursor signals of contamination, forecast hazard emergence, and quantify risk propagation across product lines or supplier networks.

Methods including multivariate anomaly detection, time-series forecasting, and Bayesian inference can predict contamination likelihood from deviations in ingredient quality, equipment conditions, or environmental monitoring signals. Similar to infectious-disease forecasting and climate-risk modelling, these approaches enable proactive interventions rather than post-hoc containment.

Embedded within privacy-preserving and explainable infrastructures, predictive AI enables HACCP systems to evolve into adaptive, continuously learning risk controls that reduce recall probability, enhance supply-chain robustness, and strengthen public health protection.

What technical or organisational barriers limit the deployment of predictive AI in food systems or similar contexts?

The deployment of predictive AI in food systems is primarily constrained by fragmented and unstructured data architectures. Most food safety information, including HACCP plans, monitoring results, supplier data, and environmental records, still exists as non-standard documents rather than machine-interpretable datasets, limiting model training and interoperability.

Technical infrastructure gaps persist in small and mid-sized enterprises, where secure cloud computing, sensor integration, and privacy-preserving analytics are not uniformly adopted.

Organisational barriers further hinder deployment: risk-averse cultures prioritise compliance over innovation, and the absence of agreed data standards complicates cross-company data sharing needed for multi-party risk forecasting.

Lessons from health and environmental domains show that predictive systems become viable only when governance frameworks, explainability requirements, and incentive structures are aligned to support data interoperability, secure model access, and human-centred oversight rather than ad hoc digitalisation.

What are the key enablers and barriers for data sharing in the food industry, and how can AI help?

Data sharing across food systems is often limited by concerns over commercial sensitivity, lack of standardised data models, and uneven digital maturity among actors. HACCP and supplier data, for example, are still predominantly stored as proprietary documents, making them difficult to exchange without disclosing sensitive information. Trust and interoperability therefore become prerequisites for collaboration.

Predictive AI can enable new sharing models by operating on securely federated data rather than requiring centralised access. Privacy-preserving techniques such as federated learning, differential privacy, and secure multiparty computation make it possible to train risk-forecasting models without exposing confidential information.

When combined with standardised, machine-interpretable data structures, such as those emerging from digital HACCP systems, these technologies create incentives for collaboration by allowing companies to contribute to shared risk intelligence while retaining control over their data. AI thus becomes both a technical enabler and a governance mechanism for multi-party food safety resilience.

How do you approach explainability and trust in AI systems used for risk prediction?

Explainability and trust in predictive AI for food-risk management require models to support traceable reasoning rather than opaque outputs. Because food safety control decisions, such as those guided by HACCP decision trees for critical control points (CCPs), carry regulatory implications and directly influence consumer protection, risk-forecasting systems must demonstrate how a prediction derives from underlying hazards, process parameters, or supplier evidence.

Techniques such as feature attribution (e.g. SHAP values), Bayesian reasoning, and rule-augmented machine-learning can expose the contribution of specific data sources to predicted risk levels. Human oversight remains central: AI outputs should be presented as decision support with uncertainty estimates, auditable provenance, and clear links to recommended control actions.

As demonstrated in other industries, interpretable risk scores, transparent audit trails, and expert validation build trust among regulators and operators. In alignment with EU AI Act requirements, trustworthy AI must pair technical transparency with documented governance protocols, ensuring that modelling reinforces (rather than replaces) scientific judgement in food safety.

What considerations guide the design of sustainable AI infrastructure in your organisation or field?

Designing sustainable AI infrastructure for food risk prediction requires balancing model performance with computational efficiency and environmental impact. In food systems, where risk models may run continuously across distributed facilities, lightweight architectures can outperform large energy-intensive models by focusing on structured domain data, such as standardized hazards, CCP logic, or supplier risk attributes, rather than broad,

Position Statement from DQFS

unbounded learning. Model compression, edge deployment for on-site processing, and selective retraining strategies reduce cloud usage and energy demand.

Sustainability also depends on governance: transparent data standards help minimise redundant computation, and privacy-preserving methods (e.g., federated learning) reduce the need to centralise large datasets, lowering storage and transfer costs. Lessons from climate informatics show that targeted, domain-specific models consistently produce more stable and resource-efficient predictions.

Applied to food safety, sustainable AI infrastructures prioritise responsible scaling, domain-guided modelling, and lifecycle monitoring of model performance and compute, ensuring that resilience gains do not create new environmental burdens.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems or cross-sector collaboration?

By 2035, a transformative breakthrough would be the emergence of a global, privacy-preserving "Predictive Food Risk Commons" that unites regulators, industry, laboratories, and public health agencies through secure, explainable AI. Such an infrastructure would allow models to continuously learn from anonymised HACCP data, pathogen genomics, climate signals, trade flows and supply-chain disruptions, enabling probabilistic forecasting of contamination pathways and system stressors without requiring any participant to relinquish data ownership.

Federated learning, synthetic data, and causal AI would support early detection of emerging hazards and simulate how control interventions alter risk trajectories, while domain-aware reasoning systems could propose validated countermeasures and quantify uncertainty. Integrated with the EU AI Act, a food safety common would embed human-auditable risk logic, ethical stewardship, interoperability standards, and energy-efficient computation.

Success would mean food safety systems capable of anticipating threats collaboratively, safeguarding public health through shared foresight, and strengthened by collective intelligence rather than fragmented effort.

FOOD SAFETY & AGRICULTURE APPLICATIONS

HORIZON Smart Droplets Project

Position Statement from Agricultural University of Athens (AUA)

ΓΕΩΠΟΝΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
AGRICULTURAL UNIVERSITY OF ATHENS

Authors:
Professor **Spyros Fountas**
Borja Espejo-Garcia, Postdoctoral Researcher
Department of Natural Resources Development & Agricultural EngineeringOASEES Project Coordinator

Where can AI add the most value in food safety decision-making?

AI adds the most value to food safety when it integrates heterogeneous signals into early, actionable decisions that prevent hazards before they enter the food chain. This encompasses weather, phenology, pest pressure, and soil/water quality, as well as plant-level detection of diseases that typically emerge in small, rapidly expanding clusters. Because symptoms of key pathogens are detectable far earlier than conventional laboratory analyses (**Dhaka *et al.*, 2021**), AI-powered tools, such as smartphone-based detection developed in projects like NextGenBioPest, dramatically strengthen early warning capabilities.

Moreover, these plant-level insights become even more powerful when integrated into system-level decision workflows. For instance, in the Smart Droplets project, they are fed into prescription maps, adaptive Direct Injection System (DIS) spraying, and digital twin simulations (**Zhang et al., 2025**) that reduce chemical loads, drift, and residue-related risks while exploring "what-if" strategies for choosing the safest intervention paths.

How is AI helping detect and respond to emerging risks in your sector?

AI is transforming the detection of emerging risks in agriculture by enabling continuous, plant-level surveillance and rapid diagnosis of early symptoms that would otherwise go unnoticed. Many fungal pathogens emerge in small, localised outbreaks that expand rapidly, and their early signs are often difficult to detect using conventional methods. AI-powered vision systems, whether mounted on UAVs, retrofit tractors, or embedded in smartphones, now provide near–real-time disease detection.

Position Statement from AUA / Smart Droplets

AI also plays a central role in responding to these risks by linking detection outputs to dynamic, operations-level decisions. In Smart Droplets, for instance, these visual insights are fed directly into Digital Farm Twins, where layers of canopy health, droplet deposition, and soil moisture (among others) can be integrated to simulate how risks evolve in space and time. By forecasting chemical drift under changing environmental conditions, digital twins along with multimodal AI can serve as predictive tools that alert stakehodlers to risks affecting water bodies, pollinator habitats, or adjacent food-producing areas before incidents occur.

What are the main challenges in deploying AI for resilient food systems?

The primary challenges in deploying AI for resilient food systems stem from issues of data quality, availability, and representativeness. Most actors still face limited labelled datasets, uneven data governance, and highly heterogeneous field conditions. Models trained on one crop, region, or season often fail in another due to distribution shifts, label noise, and rapidly changing pathogen ecologies. Smartphone-based systems or precision-spraying workflows in Smart Droplets, therefore, require extensive field trials across variable canopies, climates, and imaging conditions to ensure reliability.

A second major challenge lies in building interoperable, trustworthy, and continuously updated data ecosystems that allow AI models to remain robust over time. The Common European Agricultural Data Space (CEADS) aims precisely to address today's fragmentation by enabling secure sharing of farm, industry, and public data with harmonised connectors, vocabularies, and governance, prerequisites for training models that generalise across borders and seasons (**Stefanidou et al., 2025**).

What does trustworthy AI look like in the context of food safety?

Trustworthy AI in food safety requires systems that are transparent, auditable, and scientifically interpretable throughout their entire lifecycle. Under the EU AI Act , high-risk agricultural applications, such as automated disease detection or adaptive spraying, must demonstrate rigorous risk management, traceability and human supervision while maintaining auditable data processing and defining clear human-in-the-loop checkpoints, particularly when AI decisions may affect chemical applications.

Equally essential is explainability; AI models must reveal why they make specific recommendations, in ways that align with biological reality and can be inspected by farmers, cooperatives, and regulators. Techniques such as Grad-CAM (**Selvaraju et al., 2016**) allow computer-vision systems in Smart Droplets or NextGenBioPest to highlight the specific canopy regions, lesions, or grape clusters that triggered a classification or a spraying adjustment.

How can AI support cross-border coordination and data sharing to enhance food system resilience?

AI can enhance cross-border coordination by promoting the agricultural sector toward shared data standards and infrastructures that facilitate the seamless flow of information. As AI systems for disease detection, precision spraying, and digital-twin modelling become more widespread, they expose the need for accelerating the development and adoption of the Common European Agricultural Data Space (CEADS), which is expected to provide a secure, interoperable backbone for exchanging agronomic, environmental, and regulatory data across Member States.

Position Statement from AUA / Smart Droplets

CEADS' governance frameworks ensure that heterogeneous AI systems, from smartphone disease detectors in NextGenBioPest to tractor-mounted computer-vision systems in Smart Droplets, can operate beyond local silos while preserving traceability, privacy, and data protection. Moreover, imagery and sensor data captured by drones, smartphones, and spraying platforms can be ingested into CEADS-compatible Farm Management Information Systems (FMIS) (**Fountas et al., 2015**), enriching digital twins with cross-regional context on disease pressure, canopy structure, spray deposition patterns, and soil moisture.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems?

By 2035, a breakthrough would be the deployment of multimodal digital twins that continuously ingest heterogeneous data streams (e.g., field sensors, UAV imagery, hyperspectral cameras, soil and water probes, omics assays, and pesticide residue analytics) to infer the causal links between agronomic practices and food safety, biodiversity, and yield outcomes in the long term. Unlike task-specific models, these multimodal AI systems in the context of Agriculture 5.0 (**Fountas et al., 2024**) and foundation models (**Espejo-Garcia et al., 2025**) would align spatial, visual, temporal, and biological modalities through shared embeddings, enabling robust cross-domain reasoning. This way, digital twins would act as cognitive ecosystems, simulating alternative Integrated Pest Management (IPM) strategies while dynamically quantifying trade-offs among efficacy, environmental persistence, carbon footprint, and ecological impact. Moreover, they would generate auditable, explainable reports that conform to the EU AI Act, ensuring that every prescription or simulation is traceable, justifiable, and reproducible.

- *Dhaka, V.S., Meena, S.V., Rani, G., Sinwar, D., Kavita, K., Ijaz, M.F., & Woźniak, M. (2021). A Survey of Deep Convolutional Neural Networks Applied for Prediction of Plant Leaf Diseases. Sensors (Basel, Switzerland), 21.*
- *Espejo-García, B., Güldenring, R., Nalpantidis, L., & Fountas, S. (2025). Foundation vision models in agriculture: DINOv2, LoRA and knowledge distillation for disease and weed identification. Comput. Electron. Agric., 239, 110900.*
- *Fountas, S., Carli, G., Sørensen, C.A., Tsiropoulos, Z., Cavalaris, C., Vatsanidou, A., Liakos, B., Canavari, M., Wiebensohn, J., & Tisserye, B. (2015). Farm management information systems: Current situation and future perspectives. Comput. Electron. Agric., 115, 40-50.*
- *Fountas, S., Espejo-García, B., Kasimati, A., Gemtou, M., Panoutsopoulos, H., & Anastasiou, E. (2024). Agriculture 5.0: Cutting-Edge Technologies, Trends, and Challenges. IT Professional, 26, 40-47.*
- *Selvaraju, R.R., Das, A., Vedantam, R., Cogswell, M., Parikh, D., & Batra, D. (2016). Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. International Journal of Computer Vision, 128, 336 - 359.*
- *Stefanidou, A., Radoglou-Grammatikis, P.I., Argyriou, V., Sarigiannidis, P.G., Varlamis, I., & Papadopoulos, G.T. (2025). State of Play and Future Directions in Industrial Computer Vision AI Standards. 2025 IEEE Conference on Artificial Intelligence (CAI), 1528-1533.*
- *Zhang, R., Zhu, H., Chang, Q., & Mao, Q. (2025). A Comprehensive Review of Digital Twins Technology in Agriculture. Agriculture.*

CROSS-DISCIPLINARY
SCIENCE & TECHNOLOGY

CROSS-DISCIPLINARY SCIENCE & TECHNOLOGY

Kiel University

Kiel University
Christian-Albrechts-Universität zu Kiel

Authors:
Professor **Dr. Martina Gerken**
Dr. Christian Linke, Researcher
Dr. Igor Titov, Researcher
Faculty of Engineering
Nicholas Enders, Researcher
Institute for Innovation Research

How can predictive AI technologies strengthen resilience in food systems or related domains?

Predictive AI enables smarter decisions in farming by forecasting crop needs - such as fertilisation, irrigation, and plant protection - based on soil, weather, and satellite data. This reduces input waste, improves yield stability, and lowers environmental impact. It also enhances food system resilience by reducing reliance on non-renewable resources and minimizing exposure to climate variability.

AI can forecast pest outbreaks or disease risks from weather and historical patterns, stabilizing production and limiting price shocks. To be effective, all such models rely on timely, site-specific field data. New sensor technologies under development, such as automated in-soil lab-on-a-chip systems, can address this data gap by providing continuous, high-quality nutrient measurements, enabling real-time decisions and better AI model performance.

What technical or organisational barriers limit the deployment of predictive AI in food systems or similar contexts?

A major barrier is the lack of suitable sensors for key factors such as plant-available nutrients, soil biology, and plant health. Manual sampling and laboratory testing are labour-intensive and infrequent, resulting in sparse datasets with limited resolution. Harsh outdoor conditions often cause data gaps and equipment failures.

Additionally, there is poor interoperability across fragmented IT systems, each using different formats and classifications. The complexity of biological systems - with non-linear, time-delayed, and adaptive responses - makes modelling challenging, especially with limited data. AI tools also require processing power and connectivity,

Position Statement from Kiel University

which are often unavailable in rural areas. Finally, digital systems can be costly, and the value they generate may not directly benefit the data provider, limiting adoption.

What are the key enablers and barriers for data sharing in the food industry, and how can AI help?

Data sharing is hindered by technical, economic, and organisational barriers. Heterogeneous data models and formats across IT systems othen result in loss of meaning during transfer. Large players benefit from vendor lock-in and lack incentive to open up. Farm-level data are costly to collect and may benefit downstream actors more than farmers, discouraging participation. AI could play a role by translating between data formats and lowering integration costs, but only if underlying models are accessible. Privacy regulations such as the GDPR also restrict the use and transfer of personally identifiable data. A trusted governing body is required to harmonise data standards and ensure fair, secure use across the value chain, balancing costs and benefits to all stakeholders.

How do you approach explainability and trust in AI systems used for risk prediction?

Transparency in AI models is crucial for building user trust, particularly in agriculture where decisions directly affect livelihoods. A key strategy is using verifiable data sources. At Kiel University, we are developing a buried lab-on-a-chip system that generates time-series data on plant-available nutrients, enabling AI models to be trained on trustworthy ground-truth data. Beyond data quality, explainability also depends on model design and communication. Models should include mechanisms to visualise trends, identify outliers, and demonstrate how input changes affect outputs. Highlighting year-on-year trends or comparing sensor locations can help users interpret the system and foster confidence in decision support tools - especially when users maintain oversight of the outcomes.

What considerations guide the design of sustainable AI infrastructure in your organisation or field?

Sustainability in digital agriculture involves balancing functionality, energy use, and long-term accessibility. AI models should be efficient and tailored to their operational environment. In the field, limited power and connectivity make lightweight models and edge computing critical. At Kiel University, we design sensor platforms to be low-maintenance, buried beneath the surface, and compatible with field machinery. Our infrastructure choices also consider the long-term total cost of ownership. On the software side, we prioritise interpretable models and visual tools that help users understand outcomes. Sustainability also includes data ethics and openness: reusable models and open APIs can help avoid redundant development and promote broader system integration.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems or cross-sector collaboration?

By 2035, we envision widespread use of real-time sensor networks combined with remote sensing and AI to deliver precision recommendations tailored to specific crops, soils, and climates. Affordable, autonomous sensor technologies will make it economically viable to collect the dense, high-quality data AI models require. AI-based translation systems will improve interoperability across platforms, actors, and borders. We also anticipate more public infrastructure, such as government-backed cloud services, to reduce dependency on private providers and ensure compliance with data protection regulations. Ideally, breakthroughs in explainable AI will make complex models understandable to all users, bridging the gap between decision support and farmer intuition in an increasingly digital food system.

CROSS-DISCIPLINARY SCIENCE & TECHNOLOGY

Foundation for Research and Technology - Hellas (FORTH)

Author:
Stavros Pissadakis, Director of Research
Institute of Electronic Structure and Laser

How can predictive AI technologies strengthen resilience in food systems or related domains?

Optical sensing technologies rapidly penetrating into the Precision Agriculture (PA) domain, offering a wealth of new types of data, that were not previously available to farmers, agronomists and researchers. Predictive AI technologies are now envisioned (and in some cases already applied) to interpret and translate these optical sensing data into simple and perceivable parameters for immediate and reliable monitoring of the crops and the field.

The multi-parametric sensing and risk evaluation processes may include early detection of pest threats, adjustment of fertilisation in terms of macro- or micro-nutrients, or monitoring of abiotic stresses of plants and trees. Such multi-parametric data analyses powered by computational and machine learning tools, shift Agriculture from the traditional empirical mode into the Precision "AI-powered decision support and sensor fusion" mode, saving natural and cultivation resources, also reducing major cultivation risks, which can downgrade both the quality and quantity of crop yield.

The principal idea will be to use optical sensing technologies together with Predictive AI models not for formulating crucial decisions, but for assisting the end users in reducing the uncertainty of the decisions taken.

What technical or organisational barriers limit the deployment of predictive AI in food systems or similar contexts?

A major issue hurdling the wider deployment of Predictive AI in optical sensing subsystems and methods employed in the modern agricultural sector is the lack of standardisation covering both hardware and data management systems. This lack of standardisation is related to the rapid penetration of several new optical sensing technologies into the Precision Agriculture field, also with the difficulty of translating and correlating crucial agricultural parameters (i.e. pesticide residue levels) with specific optical quantities.

The fusion of data obtained from different optical sensing platforms and protocols into unified risk prediction toolkits using Prediction AI models still remains a challenge, since the type and particular characteristics (i.e. temporal distribution, accuracy, etc.) of the data available, may lead to conclusions that are predominantly based on the interpretation method, yet based less on the measurements obtained on the physical (sensor) layer.

What are the key enablers and barriers for data sharing in the food industry, and how can AI help?

Optical sensing technologies, especially those based on fibre-optic sensors, provide point measurements in agricultural fields, requiring deployment over a grid for accumulating substantial sampling for a single user. It is obvious that the monitoring of a wider agricultural area, covering for instance a region assigned as a Protected Designation of Origin, requires an extended network of sensing points, which may be distributed along different users; there data sharing is inevitable for creating necessary databases.

Privacy-preserving AI, including blockchain protocols, can support the sharing of specific types of data between end users and network moderators (i.e cooperatives), without compromising privacy rules or disclosing sensitive cultivation practices. The same Privacy-preserving AI can be employed for diluting and integrating the data available into a greater model of higher validity, being broadly available so interested end users can access it to cross-check the measurements of their sensing points and consolidate results.

How do you approach explainability and trust in AI systems used for risk prediction?

A catalyst that will boost the adoption of optical sensing technologies into the Precision Agriculture by a large number of farmers and agronomists will be the establishment of trust and transparency across both the optical layer and the AI signal processing technologies used. The optimum way to enhance trust of among end users of these emerging technologies is through their association with and validation against realistic, success cases.

Therefore, these new hardware and AI analysis technologies should be thoroughly validated in real field cultivation cases, involving a human-in-the-loop mode, and showcasing tangible, measurable benefits for the main stakeholders. The credibility of those test cases and the relevant risk assessment performed using AI models may be undertaken by independent organisations, for example, international research centres or large scale cooperatives, while involving domain experts with sound knowledge in photonics, AI and agronomy.

Agricultural resources (i.e. equipment, fertilisers or pesticides) and other high tech (i.e. photonic components manufacturers) may also be involved as beneficiary players while populating a broader value chain.

What considerations guide the design of sustainable AI infrastructure in your organisation or field?

Precision Agriculture deploys along three main sustainability priorities: efficient use of resources, improvement of final agricultural product quality for consumers, and cultivation with minimal environmental impact. These priorities can be further implemented into the framework of circular economy, tailored to each crop, area, agricultural practices and business model.

Optical sensing technologies can readily serve all these three sustainability priorities, however, the fusion of the optical sensing data into a sustainability-oriented cultivation and business model is not yet straightforward. The key challenge relates to the profit margin generated for the farmer either operating in a personal or cooperative

Position Statement from FORTH

business model, without compromising sustainability priorities, and without significantly disturbing other parts of the value chain (i.e. final consumer or agrochemical companies). Thus, a more generalised profit-sustainability AI model should be implemented, seamlessly fitting to the hardware technologies and business models available.

Looking ahead to 2035, what breakthrough would you like to see in AI for resilient food systems or cross-sector collaboration?

Precision Agriculture comes to change the way we cultivate, produce, certify, transport and sell crops. Sensing technologies, including those of optical sensing, are cornerstone components into PA, with the sensing data obtained dominating most parts of the extended and cross-sectorial value chain of Agro-food. Future AI models should operate for generating win-win opportunities for all players involved in this value-chain, especially for farmers that are exposed to frequent and unpredictable cultivation and financial risks.

A generalised AI model should cover the whole ecosystem of Agrofood sector, bringing closer all players involved, while providing transparency of operations and transactions, allowing trustworthy tracking of the products from Farm to Fork. Such a comprehensive AI model could secure crop quality, ensure fair profit for all player involved, sustainability practices and timely availability of the products to the markets, while avoiding disruptions, which can lead to scarcity of products, and price volatility.

Contributions by Organisations & Initiatives

ΓΕΩΠΟΝΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
AGRICULTURAL UNIVERSITY OF ATHENS

BioSim
Biomedical Simulations and Imaging Laboratory

CERTH
CENTRE FOR RESEARCH & TECHNOLOGY HELLAS

Consiglio Nazionale delle Ricerche
ISTC
CNR - Istituto di Scienze e Tecnologie della Cognizione

Met Office

DEMOKRITOS

FORTH
FOUNDATION FOR RESEARCH AND TECHNOLOGY - HELLAS

CAU
Kiel University
Christian-Albrechts-Universität zu Kiel

UCD DUBLIN

Smart Droplets

PLIADES

HACID

OASEES

ROBOSAPIENS

EFF-CoP
European Food Fraud Community of Practice

Agroknow

FOODAKAI

EFRA
EXTREME FOOD RISK ANALYTICS

EFRA
EXTREME FOOD RISK ANALYTICS

Let’s Ensure Safer Food For All

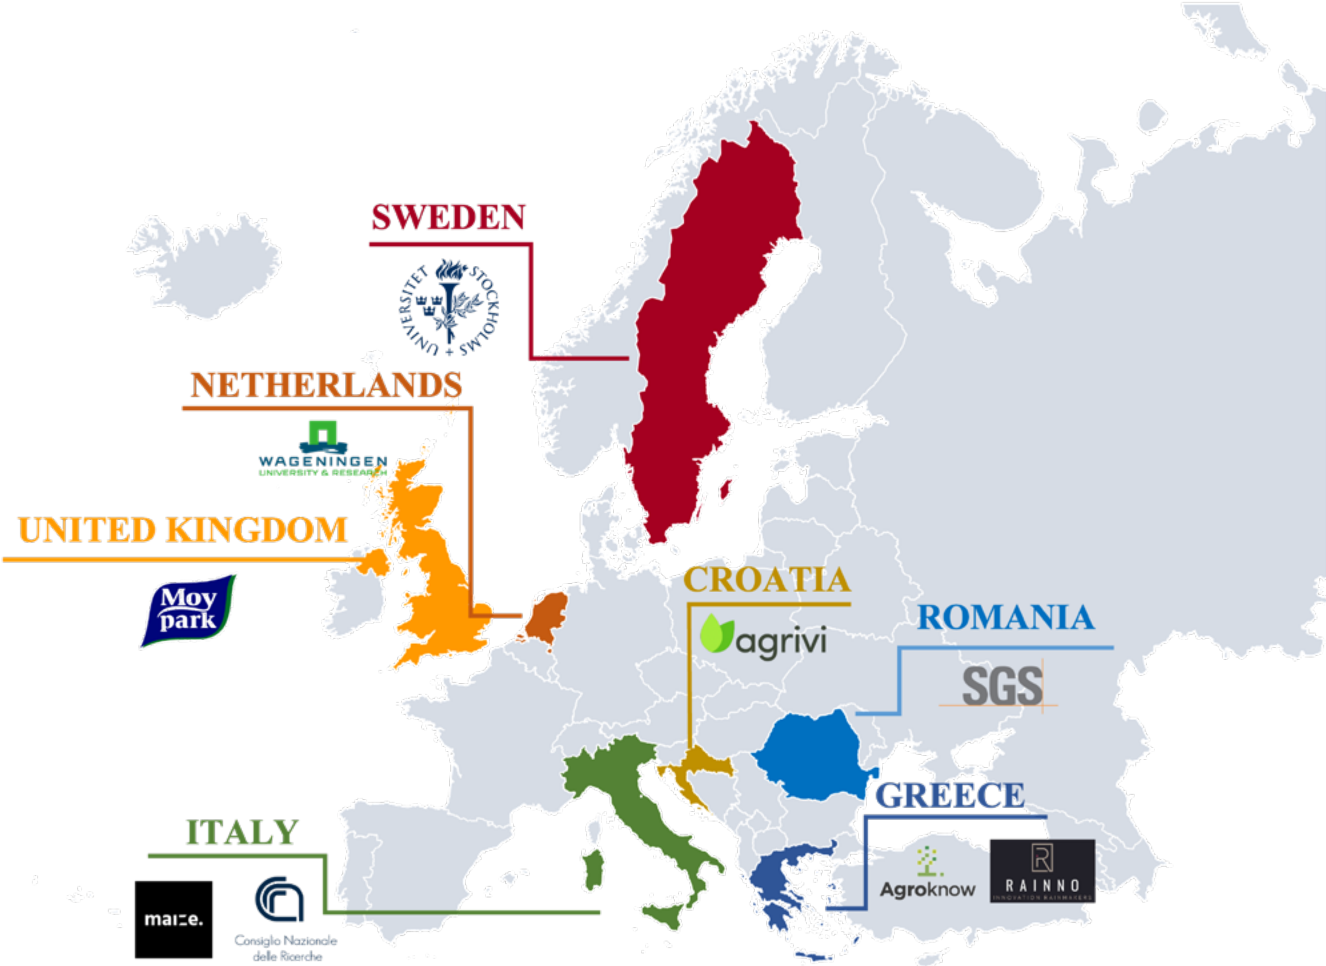

Contact Us: www.efraproject.eu | www.agroknow.com

/EFRA

@EfraProject

Disclaimer:
Funding for this research has been provided by the European Union’s Horizon Europe research and innovation programme EFRA (Grant Agreement Number 101093026). Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission-EU. Neither the European Union nor the granting authority can be held responsible for them.