Public Document

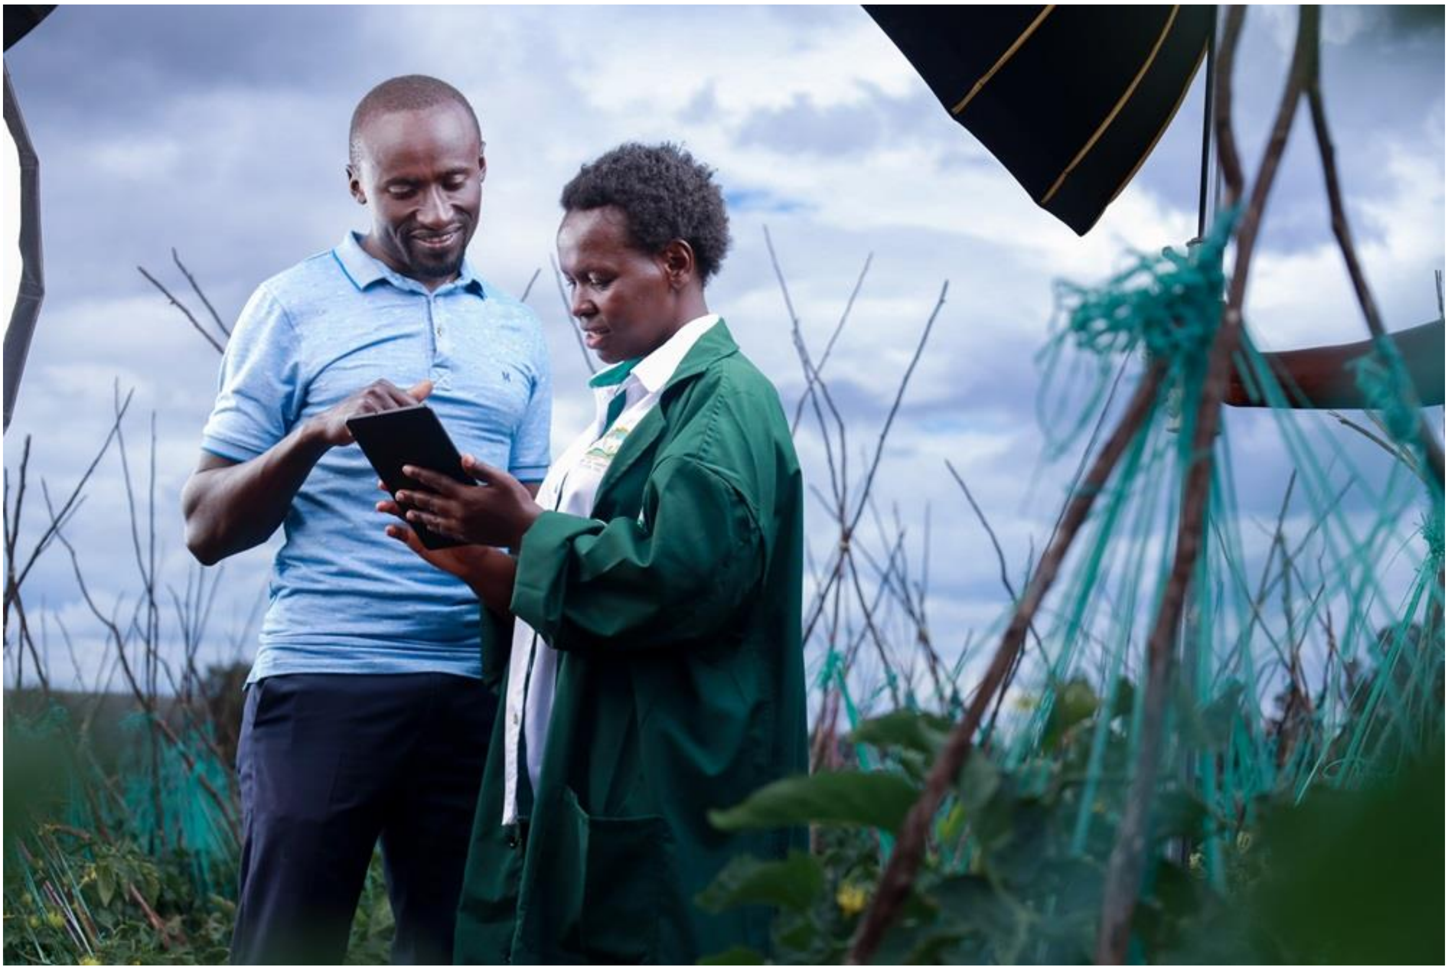

GenAI for Agriculture Advisory (GAIA)

Risk and Ethics Assessment and the GAIA Ethics Toolkit

CABI

November 2025

KNOWLEDGE FOR LIFE

Risk and Ethics Assessment and the GAIA Ethics Toolkit

The writing of this report was carried out during the second phase of the Generative AI for Agriculture Advisory (GAIA) project, led by IFPRI (CGIAR). CABI gratefully acknowledges the funding provided to this project by the Gates Foundation and the Foreign, Commonwealth and Development Office, UK (FCDO).

Table of Contents

Executive Summary

Generative AI (GenAI) is rapidly reshaping agricultural advisory systems, offering new opportunities for climate resilience, productivity gains, and improved decision-making. Yet its adoption in agriculture—particularly in smallholder contexts—comes with significant socio-technical and ethical risks. These risks arise from data quality limitations, unequal power dynamics, opaque algorithms, and gaps in accountability and governance. This report outlines these emerging challenges and presents the rationale, structure, and expected impact of a forthcoming GAIA Ethics Toolkit designed to ensure responsible development and deployment of GenAI-enabled agricultural advisory systems.

Global conversations around data governance, including outcomes from the 2024 UN World Data Forum, stress the importance of equitable and FAIR data ecosystems. In agriculture, where digital innovation intersects closely with farmer livelihoods, environmental stability, and national strategies, responsible AI development is essential. Applications such as crop management chatbots, long-lead weather forecasting, and AI-driven pest detection tools already demonstrate strong potential. However, they also highlight the vulnerabilities introduced by poor data representativeness, model opacity, privacy risks, and the marginalisation of traditional knowledge.

Key ethical risks include:

- Accuracy and reliability issues driven by biased, incomplete, or unverified datasets.
- Undermining of indigenous knowledge when AI systems privilege formal scientific data.
- Lack of explainability in systems like crop recommendation engines, which erodes user trust.
- Accountability gaps due to unclear responsibility for data provenance, model decisions, and harmful outcomes.
- Privacy and security concerns, especially where farmers lack bargaining power over data use.
- Data monopolisation, potentially deepening inequities between global agribusinesses and smallholders.
- Environmental impacts of energy-intensive AI models.
- Misalignment between technological optimisation and farmer needs, with low adoption rates in many low-income regions.

To address these challenges, the GAIA Ethics Toolkit will provide a practical governance resource to guide developers, researchers, and implementing partners. Organised around the AI/ML lifecycle—from defining the use case to post-deployment monitoring—the toolkit will offer checklists and prompts for identifying risks, ensuring legal compliance, establishing feedback and redress mechanisms, and promoting iterative improvement. User stories and stakeholder engagement will shape the toolkit to ensure it is accessible, actionable, and tailored to real-world agricultural advisory contexts.

When adopted, the toolkit is expected to help developers embed ethics-by-design practices, improve model transparency and reliability, strengthen user trust, and support responsible scaling of GenAI-enabled advisory services. Over time, this foundation can promote more equitable, sustainable, and evidence-based agricultural systems that improve farmer livelihoods and climate resilience.

The coming phase of work under GAIA will focus on toolkit development, iterative testing, and validation with GenAI developers and partner organisations. A functional version is expected by mid-2026, with refinement and a full launch planned for early 2027.

Introduction

The fifth UN World Data Forum (November 2024) underscored the importance of equitable data ecosystems in advancing the Sustainable Development Goals (SDGs). Within this context, Area 2 of the Medellín Framework highlights the need for Findable, Accessible, Interoperable, and Reusable (FAIR) data, improved communication strategies for data-driven decision-making, and enhanced data literacy. As artificial intelligence (AI) — particularly Large Language Models (LLMs) — continues to shape digital transformation, data access remains central to development outcomes. Recognising both the promise and the challenges of operationalising the FAIR principles introduced in 2016, the Gates Foundation and CABI developed the FAIR Process Framework (FPF) to guide agricultural development investments toward more systematic and inclusive adoption.

Agriculture, a cornerstone of the SDGs and national AI strategies such as those of India, stands at the intersection of digital innovation and ethical complexity. While AI-driven tools can improve efficiency, profitability, and climate resilience, they also raise concerns over data monopolies, potential to drive reduced crop diversity, and farmer autonomy. Many of these risks are rooted in data quality, security, and governance. This brief examines these data-centric risks in generative AI use cases for ag advisory, and how we intend to address them through the development of an ethics toolkit.

Applications of GenAI in Ag advisory

Ongoing research and development efforts are exploring various AI applications in agriculture and food systems. Here we highlight some of the prominent applications of GenAI systems to improve ag advisory delivery and scale:

i. Crop management chatbots: AI tools, such as chatbots, can extract insights from diverse datasets to provide guidance on crop management. For example, Digital Green's multilingual AI-powered chatbot serves as both a content retrieval and delivery tool. It allows Agricultural Extension Agents (EAs) to send and receive customised messages and query the database for specific information on issues they encounter in the field, tailored to local conditions. Additionally, data dashboards and a farm-level report generator help EAs assess farmers' needs and plan their activities. These tools also enable extension system managers to review reports on EA and service provider activities, as well as evaluate the effectiveness of EAs and the content they provide.[1]

ii. Weather forecasting tools with better lead time: The system *"Sub-Seasonal to Seasonal Forecasting for Agriculture in Bangladesh"* was launched to help farmers manage their production plans. Farmers typically receive weather forecasts that extend from daily updates up to five days ahead. However, this short notice often leaves them with insufficient time to take necessary actions to protect their crops from unfavourable weather. With the new automated system, farmers receive weather forecasts for the coming four weeks to three months, enabling them to make better production decisions based on anticipated weather patterns.[2]

---

[1] Digital Green, 'FarmerCHAT: Enabling Farmer-Centric Advisory Services' (*Digital Green*) https://digitalgreen.org/farmerchat/ accessed 28 November 2025.
[2] Dhaka Tribune, 'New Weather Forecasting System Launched for Bangladesh Farmers' (*Dhaka Tribune*, 27 March 2023) https://www.dhakatribune.com/bangladesh/307718/new-weather-forecasting-system-launched-for accessed 28 November 2025.

iii. Pest detection applications: AI can help analyse vast amounts of data to support farmers in real-time decision-making. Plantix is a mobile application that uses AI to quickly identify common pests, diseases and nutrient deficiencies in a range of different crops, enabling farmers and extension workers to respond promptly. Users can upload a photo of an affected crop, which the app analyses for symptoms using image-recognition technology, leveraging a dataset of over 120 million images. The app then provides a diagnosis and suggests actions to address the identified issues.[3]

Key Ethical Risks & Considerations

Given technical challenges, digital divides, and power differentials between farmers and other market players such as agribusinesses and insurance companies, the rapid adoption of GenAI in ag advisory introduces significant risks. A primary cause of these can be attributed to poor quality, incomplete or scientifically suspect datasets. Risks can also emerge from the GenAI model (e.g., outputs are not interpretable, or prejudiced assumptions have been made during development of the algorithm), from application design (e.g., needs of women farmers are excluded) or from use (e.g., environmental impacts of training and using LLMs).

Some of these social and ethical considerations include:

i. Uncertainty in accuracy and reliability: AI systems can generate inaccurate or unreliable results, leading to significant negative impacts. This phenomenon is termed as hallucinations indicating the failure of an AI system to generate reliable and accurate outputs. For example, if an AI-powered farm advisory provides incorrect recommendations on pesticide use, it could result in crop failure, harm to the farmer and/or harm to the environment. This, in turn, can threaten food security for larger populations and cause financial losses for farmers.

   Data issues are a primary cause of AI inaccuracies. Problems can arise from biases introduced during data collection, such as sampling bias (e.g., data is only collected from specific regions), exclusion (e.g. data excludes certain demographic groups) and selection bias (e.g. dataset focus on high-value crops).[4]

   Accurate AI outputs require high-quality, comprehensive, and up-to-date data, but meeting this need presents logistical challenges. Variability in agro-climatic zones demands frequent and widespread data collection. Although low-earth orbit satellites have improved the availability of geospatial data in recent years, satellite data alone is insufficient - it must be cross-verified with ground-truth data. This process involves deploying IoT sensors and manual labour, both of which are resource intensive. Furthermore, satellite coverage remains incomplete in many regions, leaving

[3] GSMA, 18 February 2025. Detecting and managing crop pests and diseases with AI: Insights from Plantix. https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/blog/detecting-and-managing-crop-pests-and-diseases-with-ai-insights-from-plantix/ accessed 27 November 2025.

[4] Kevin Mallinger & Ricardo Baeza-Yates, Responsible AI in Farming: A Multi-Criteria Framework for Sustainable Technology Design (2024) 14(1) *Applied Sciences* https://www.mdpi.com/2076-3417/14/1/437, accessed 20 November 2025.

significant data gaps. At times, the necessary data simply isn't available, limiting the effectiveness of AI systems.[5]

ii. Undermining of traditional knowledge: AI's emphasis on scientific knowledge over traditional or indigenous practices can threaten local knowledge systems. When data collection prioritises standardised, scientific information, it risks marginalising valuable indigenous insights that have supported agricultural practices for generations. This potential loss of traditional knowledge adds another layer of distrust in AI adoption.

iii. Lack of explainability: AI systems are black boxes. For example, Crop Recommendation Systems (CRS) are digital tools designed to assist farmers in selecting the most suitable crops by analysing variables such as soil characteristics, climate conditions, and past crop performance. However, their internal workings are opaque, and the reasoning behind their recommendations is not easily understandable.[6] The opacity of CRS models, as with most AI algorithms, makes it difficult for farmers to understand how input data (such as soil pH or rainfall patterns) translates into specific crop recommendations. If farmers cannot see or comprehend the underlying logic, they will question the reliability and fairness of the outputs.

    To compound matters, if datasets are incomplete, biased, or outdated, it can skew recommendations. The opaque nature of data pipelines, from collection to processing, can make it difficult to trace how decisions are generated. This lack of transparency complicates efforts to determine whether a failure is due to incorrect data, flawed algorithms, or operational issues on the farm. Without clear explanations of how data informs the results, farmers may hesitate to adopt the technology or rely on its outputs, limiting the potential benefits of these systems.

iv. Gaps in accountability: Without modalities to trace data provenance and visibility into how data drives outcomes, both farmers and developers may pass responsibility back and forth, creating an accountability gap. Furthermore, the ambiguity of agreements and legal frameworks around data collection, processing, and sharing may result in uncertainty in who manages data and how.[7] The lack of clear accountability creates significant risks for farmers when AI systems produce harmful or incorrect outcomes. In many cases, there is no defined oversight structure to determine who should be held responsible when errors occur. The current pressure to rapidly develop and scale these technologies often prioritises speed and market size over reliability and quality, leaving little room for thoughtful ethical standards or robust adoption processes.[8]

    This challenge becomes particularly acute with autonomous farm equipment and AI-based decision support systems, where responsibility for errors—such as equipment

[5] Digital Futures Lab, AI and Climate Action in Asia: An Overview of Emerging Opportunity Areas and Socio-technical Challenges (2023), https://www.climateai.asia/reports/DFL_AI_LandscapeReport.pdf, accessed 28 November 2025.

[6] Mahmoud Y. Shams et al, Enhancing crop recommendation systems with explainable artificial intelligence: a study on agricultural decision-making (2024) 36(15) *Neural Computing and Applications* https://link.springer.com/article/10.1007/s00521-023-09391-2, accessed 28 November 2025.

[7] Rozita Dara et al, Recommendations for ethical and responsible use of artificial intelligence in digital agriculture (2022) 5 *Frontiers in Artificial Intelligence* https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.884192/full, accessed 28 November 2025.

[8] Carrie S. Alexander et al, Who is responsible for 'responsible AI'?: Navigating challenges to build trust in AI agriculture and food system technology (2024) 25 *Precision Agriculture* 146-185 https://link.springer.com/article/10.1007/s11119-023-10063-3?fromPaywallRec=false, accessed 28 November 2025.

malfunctions or inaccurate predictions—can be ambiguous. For example, AI systems designed to optimise fungicide use on crops aim to increase yields and reduce costs by recommending the ideal timing and dosage.[9] However, while the system suggests the correct application schedule, the outcome also depends on the farmer's ability to follow the recommendations precisely. If farmers suffer financial losses or reputational damage due to fungicide residue at harvest, they may struggle to find anyone accountable, even if they adhered to the AI's guidance.[10]

v. Privacy and security concerns: Farmers often lack control over how their data is collected and managed, which raises concerns about the misuse of sensitive information. If data is improperly managed, it can lead to breaches of confidentiality and a loss of farmer autonomy. Since farmland is usually privately owned, farmers may be reluctant to share data due to privacy concerns and the fear of commercial exploitation.[11] Privacy risks could include identification, reputation loss or unauthorised access to data.

Integrating data from various devices on the farm can expose farmers to risks, as third parties might gain access to business-sensitive information that they prefer to keep private. This includes details about yield quality and quantity, the inputs and pesticides used, and trade secrets related to their production methods.[12] Research also shows that data is shared with third parties without consent from farmers. This is often due to the power imbalance and information asymmetry between technology providers and farmers, resulting in the latter not being able to negotiate stronger terms when it comes to their data use, sharing and reuse.[13]

vi. Risk of data and tech monopolisation: Agribusinesses have long employed various strategies to dominate the market for agricultural inputs and machinery. They have built exclusive relationships with farmers, creating lock-ins and dependencies within their economic and technical ecosystems, limiting farmers' choices and fostering reliance on their products and services.[14] For many agribusinesses, big data has become central to their business practices - Bayer-Monsanto has been building up one of the largest agricultural databases in the world, and in 2014, Monsanto[15] acquired the weather analytics company The Climate Corporation. Big agricultural data is an intangible resource, an income-generating asset for those who control quantities of data and can further process it.

Combined with limited access to technology and digital infrastructure in global south contexts, this could widen the gap between resource-rich commercial farmers and

---

[9] Rozita Dara et al, Recommendations for ethical and responsible use of artificial intelligence in digital agriculture (2022) 5 *Frontiers in Artificial Intelligence* https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.884192/full, accessed 28 November 2025.

[10] Maaz Gardezi et al (2023), Artificial intelligence in farming: Challenges and opportunities for building trust, https://acsess.onlinelibrary.wiley.com/doi/full/10.1002/agj2.21353 accessed 28 November 2025

[11] Hongyan Zhu et al,Harnessing Large Vision and Language Models in Agriculture: A Review (2024), *Arxiv,* https://arxiv.org/pdf/2407.19679 accessed 28 November 2025.

[12] Simone van der Burg et al, Ethical aspects of AI robots for agri-food; a relational approach based on four case studies (2022) 39 AI & Society 541-555 https://link.springer.com/article/10.1007/s00146-022-01429-8 accessed 28 November 2025.

[13] Supra n. 11

[14] Sarah Hackfort et al), Harvesting value: Corporate strategies of data assetization in agriculture and their socio-ecological implications (2024) 11(1) Big Data & Society https://journals.sagepub.com/doi/10.1177/20539517241234279?icid=int.sj-full-text.similar-articles.1 accessed 28 November 2025.

[15] In 2018, Monsanto was acquired by Bayer, and is now known as Bayer-Monsanto.

resource-poor smallholders, deepening inequality. Moreover, with data often processed by international corporations, there are concerns about data sovereignty—the right of countries and communities to control how their data is collected, used, and monetized. Without adequate policies, valuable local agricultural data could be extracted without reasonable compensation.

vii. Negative environmental Impact: While AI is expected to enhance sustainability, the training and deployment of AI models, especially the Large Language Models (LLM) carry their own environmental costs. The high computational cost and environmental impact of training large models raise complex ethical concerns in particular on using AI for building climate resilience in agriculture. Training models with parameters in the billions requires large quantities of electricity which lead to CO_2 emissions. Moreover, reliance on AI for intensive farming practices may lead to unsustainable land use and environmental degradation, with long-term consequences for biodiversity, soil health, and water resources.[16]

viii. Ambiguity in suitability and end-beneficiaries: Beyond the specific challenges of deploying AI in agriculture, there are broader systemic concerns about what these technological interventions are optimising for and who ultimately benefits. A key tension lies between using technology to address climate change mitigation and adaptation versus focusing on improving yields in ways that align with farmers' best interests. For instance, while AI tools might enhance yield efficiency and reduce energy consumption, they may not necessarily help farmers secure better market prices for their crops.[17]

Moreover, despite many AI-based experiments targeting farmers, the current low adoption rates of these interventions, particularly in low-income countries, highlight the need for a deeper examination of the suitability of these technologies.

[16] Elenita Dano, 'A commentary on agricultural digitalisation for climate action in the Philippines (2023) *Digital Futures Lab* https://www.climateai.asia/reports/DFL_AI_Expert_Brief_The_Philippines_vF.pdf accessed 20 December 2024.

[17] Digital Futures Lab, AI and Climate Action in Asia: An Overview of Emerging Opportunity Areas and Socio-technical Challenges (2023) https://www.climateai.asia/reports/DFL_AI_LandscapeReport.pdf accessed 28 November 2025.

GAIA Ethics Toolkit

About the Toolkit

The **GAIA Ethics Toolkit** is envisioned as a governance and AI ethics resource tailored to generative AI applications in agriculture, specifically agricultural advisory. Its purpose is to support the GAIA program, new CGIAR Science Programs, the CGIAR Digital Transformation Accelerator and other funded investments in AI as they navigate emerging challenges around bias, transparency, accountability, and socio-technical risks in AI-enabled ag advisory systems. The toolkit is geared towards developers of such systems, both commercial and non-commercial, and aims to deliver on how AI-enabled ag advisory systems can be made more trustworthy, thereby improving their adoption and scalability within agricultural communities and ecosystems.

Central to the toolkit is a structured approach for identifying socio-technical risks and measuring potential harms associated with specific agricultural AI use cases. It will include mitigation strategies, community-engagement methods, and guidance on navigating relevant legal and regulatory frameworks. Mechanisms for harm redress and accountability will be integrated, recognizing their overlap with compliance processes. The toolkit will also link these components to different stages of the AI/ML lifecycle.

To ensure accessibility and practical adoption, the toolkit will include clear user entry points and will be structured as downloadable checklists and prompts. The toolkit is being devised recognising that similar resources on ethical approaches to AI development have been designed, but there is a paucity of such actionable tools in the agricultural sector generally, and specifically with respect to ag advisory use cases.

Finally, the toolkit's development will rely on robust user stories and engagement with stakeholders across the AI lifecycle. These narratives will surface expectations, value propositions, and potential barriers faced by different user groups—for example, helping smallholder farmers understand how AI-driven recommendations are generated, or supporting developers in implementing fairness, privacy, and safety safeguards. Through iterative testing and refinement, the toolkit will become a practical, actionable, and trusted resource for building responsible, equitable, and transparent AI systems in global agriculture.

Structure of the Toolkit

As mentioned above, the toolkit is to be structured into checklists and prompts that trigger GenAI developers to consider ethical and socio-technical flaws in their potential AI systems and address them in a timely manner. At the least, even if resolution of all risks is not possible, it will ensure a more ethics-by-design approach to the development and deployment of GenAI systems for ag advisory.

To ensure that the ethics-by-design process is integral, the checklists and prompts are mapped to the different stages of the AI lifecycle, from ideation and problem identification to actual deployment and post-deployment iterative evaluation. The broad stages this toolkit will cover are as follows:

i. **Defining the use case**
   This stage identifies the specific problem AI is meant to solve, the intended users,

expected benefits, and potential risks. It sets clear goals and ensures the AI solution is necessary, feasible, and aligned with user needs and ethical considerations.

ii. **Data collection and curation**
Relevant data is gathered, cleaned, labelled, and assessed for quality, representativeness, and potential biases. This stage ensures that the dataset accurately reflects real-world contexts and supports fair, reliable model development.

iii. **Model training**
Algorithms learn patterns from curated data to produce predictions or outputs. During this stage, developers select model architectures, tune parameters, and document decisions that may affect performance, fairness, and transparency.

iv. **Evaluation and risk assessment**
The model is tested using separate data and real-world scenarios to check accuracy, robustness, fairness, explainability, and potential harms. Risks—including social, environmental, and ethical impacts—are identified and mitigation strategies are developed.

v. **Deployment**
The model is integrated into real-world systems or products and made available to end users. This stage includes preparing user interfaces, setting access controls, managing rollout, and ensuring the system behaves as expected in diverse contexts.

vi. **Governance and policy regulation**
Formal rules, accountability mechanisms, and compliance requirements are applied to ensure the AI system aligns with legal frameworks, ethical guidelines, and organizational policies. This includes documentation, audits, and oversight processes.

vii. **Monitoring, evaluation, and improvement**
Once in use, the AI system is continuously monitored to detect performance issues, unintended harms, data drift, or misuse. Feedback loops support updates, retraining, or system redesign to keep the AI safe, effective, and responsive to user needs.

For each stage the toolkit will prescribe certain questions as prompts to cause deliberation and intentional action towards mitigating ethical and socio-technical risks. The questions will be covering one or more of the following themes, for each stage:

i. **Risk identification and mitigation**
Spotting potential technical, social, and ethical harms early and putting safeguards in place to prevent or reduce them.

ii. **Regulatory and legal compliance**
Ensuring AI systems follow all relevant laws, standards, and policies across the countries and contexts where they operate.

iii. **Feedback and complaints mechanisms**
Providing clear, accessible channels for users to question decisions, report harms, and seek redress.

iv. **Iterative improvement**
Continuously refining AI systems based on monitoring results, user input, and evolving best practices to keep them safe and effective.

An illustrative table of the checklist is presented below:

Phase	Question	Response	Reasoning
Design (ideation)	How has the problem been identified?	TBA	TBA
Deployment (Legal compliance)	Are there any privacy laws that affect the GenAI system?	TBA	TBA
Deployment (Feedback)	Is the feedback collected digitally? Can feedback be submitted in other (non-digital) formats?	TBA	TBA

Expected Outcomes from the use of the Toolkit

The GAIA Ethics Toolkit addresses a core challenge in agricultural advisory: farmers often lack timely, context-specific insights, while developers struggle to build reliable and explainable GenAI systems due to barriers such as limited data access and insufficient guidance on ethical design. By providing developers with routes to secure appropriately licensed content and structured tools to consider ethical risks early in the lifecycle, the toolkit equips them with practical resources to navigate design, development, deployment, and compliance challenges. Its checklists and guidance on risk mitigation, legal obligations, and feedback mechanisms create a consistent foundation for responsible GenAI development.

As developers apply these resources, they become more aware of potential risks and unintended harms and begin embedding safety-by-design practices into their work. This leads to more reliable and trustworthy GenAI models, strengthening user confidence and enabling wider adoption of AI-driven advisory services among farmers and institutions. Over time, these improvements support more sustainable, equitable, and data-driven agricultural systems, contributing to better farmer livelihoods, productivity gains, and greater climate resilience—assuming continued engagement from developers, farmers, and institutions, and stable access to data and digital infrastructure.

Conclusion and Way-Forward

While a key objective for this report was to identify socio-technical risks and ethical considerations around the specific use case of GenAI systems for improving the delivery and quality of ag advisory, the toolkit is a key deliverable from an AI governance perspective under GAIA phase II. In 2026, we will be working on the actual development of the toolkit through an iterative process, validating our understanding and compilation of questions and prompts from GenAI developers, relevant Foundation POs, and project partner organisations. We intend to have a functional version of the toolkit ready by mid-2026 and spend the remainder of the year finetuning the same before its formal launch in early 2027.