

Appendix A: Workshop Case Study Briefing Document

In this guide

[In this guide](#)

1. [Food Standards Agency Science Council Report of Project 'Artificial Intelligence Applications in Food Safety and Authenticity'](#)
2. [Authors, Acknowledgements and Declarations of Interest](#)
3. [Executive Summary](#)
4. [Recommendations to the FSA](#)
5. [Introduction](#)
6. [Methodology](#)
7. [Results and Discussion](#)
8. [Conclusions](#)
9. [Appendix A: Workshop Case Study Briefing Document](#)
10. [Appendix B: Workshop Participant List](#)
11. [Appendix C: Workshop Case Study Responses](#)
12. [References](#)

Food Standards Agency Workshop: Ensuring Safe and Trustworthy Application of AI in Food Safety and Assurance.

Preamble

As the use of artificial intelligence (AI) expands across the food sector, new opportunities and regulatory challenges are emerging. This workshop, hosted by the Food Standards Agency, brings together stakeholders from food safety, regulation, food production and supply and the AI community. Its objective is to explore how AI is, and might be, applied in the key domains of food safety and assurance. Using a series of case studies, participants will examine current and emerging applications ranging from risk assessment and certification to

document inspection and to visual detection of defects. The workshop will consider the opportunities and challenges arising from AI technologies, where regulatory frameworks may need to evolve, what assurance mechanisms are required, and where potential gaps or risks may arise. The case studies will serve as a shared reference point for structured discussion, helping to identify areas where guidance, standards, or oversight could support safe and trustworthy deployment of AI across the food system.

Food Safety Management and Regulatory Compliance

Businesses employ a range of business processes and tools to ensure that food products placed on the market comply with relevant regulations and are safe to consume. Critical steps in such tools and processes require assurances that they work as designed and are fit for purpose. Such assurances will include, for example, implementation of recognized international, national and industry standards; reference to rigorous underpinning science; results of audits against such standards where applicable; evidence of operator competence where relevant; validation data (internal and external); supplier data (including traceability); customer and consumer complaints; contaminant and routine analytical and processing line data.

There is some degree of flexibility around how the above processes are implemented and ratified, as every business and product is different. However, any significant changes to the processes or procedures resulting from the introduction of AI technologies will require verification that the overall process of safety assurance and assessment of regulatory compliance is still operating as intended, and as specified in regulations across the food system. Such safeguards should encompass changes to physical steps such as harvesting and processing as well as data handling and analysis.

Case Study 1: AI Driven Safety and Regulatory Compliance Evaluation for Manufactured Foods

Food manufacturers developing complex, multi-ingredient products must conduct detailed safety evaluations and prepare food safety management plans to ensure products placed on the market are safe and comply with legislation. Key aspects of such evaluation and management planning include, for example, chemical and microbial safety, allergenicity, ingredient safety, and labelling accuracy and compliance. These assessments draw on an extensive and ever-evolving landscape of scientific, regulatory, and product data: from surveillance and monitoring data; validation data for processes and methods; shelf-life data;

toxicological studies, historical data and published case studies; incidents data; and legal thresholds. The data are highly distributed, heterogeneous, and often unstructured. Complex supply chains and product recipes give rise to the risk of food fraud involving, for example, substitution with cheaper raw materials, or falsifying data regarding the origin or identity of the product. Effective supplier controls can lower the risk of food fraud. Databased approaches such as blockchains can protect the integrity of supplier data. In addition, numerous analytical tools are available. These include nucleic acid-based approaches such as DNA barcoding and chemical tools such as those based on spectroscopic fingerprints. Most such approaches generate large amounts of raw data that requires extensive analysis and expert interpretation before action can be taken. Some approaches are still experimental and therefore contentious.

To manage this complexity, manufacturers are likely to turn to **AI systems** to support early-stage safety, regulatory and labelling decisions. These systems may include:

- **Large Language Models (LLMs)** to process regulatory documents, scientific literature, and guidance.
- **Knowledge Graphs** and **ontologies** to map relationships between ingredients, allergens, and known risk pathways.
- **Multimodal AI** to integrate structured data (e.g., ingredient lists, concentrations, batch records) with unstructured text (e.g., literature or safety reports).
- **Rule-based** and **ML systems** to simulate or flag risks based on novel ingredient combinations or emerging science.

These tools can significantly accelerate product development while improving consistency and thoroughness. However, the use of AI for the assessment of regulatory compliance and safety evaluation raises critical assurance questions for the Food Standards Agency.

Key Questions for the workshop:

- 1. How can the FSA be assured that AI systems used for allergenicity and compositional risk assessments have accessed, interpreted, and applied the correct scientific and regulatory data across all relevant domains?**
- 2. What standards should govern the transparency, traceability, and reproducibility of AI-derived risk assessments, particularly when used to justify labelling and safety decisions?**
- 3. How do we validate that AI systems can identify emerging risks or uncommon ingredient interactions, and not just replicate existing knowledge—especially when considering public health risk?**

Case Study 2: AI-Supported Data Pack Generation for Third-Party Certification and Assurance

Certification and assurance schemes require food producers and suppliers to demonstrate compliance with a wide range of standards across the entire food supply chain from farm to fork. Audits and inspections to ensure compliance with specified standards cover all aspects of food production, including food safety, traceability, production methods, worker safety, and environmental protection. Historically and currently, audits depend on the creation of detailed “data packs,” which draw on diverse, often fragmented, datasets across multiple systems and formats.

To automate the development of these “data-packs”, software developers are likely to explore the use of **multimodal AI systems**. These systems would integrate:

- **Natural Language Processing (NLP)** to interpret unstructured text such as process and quality records such as treatments on farm and on production line quality checks in manufacturing.
- **Optical Character Recognition (OCR)** to digitise and extract information from scanned or handwritten documents.
- **Tabular and Structured Data AI** to interpret spreadsheets, XML data, and inputs from farm and food manufacturing management software.

- **Document Question Answering (DocQA) and Retrieval-Augmented Generation (RAG)**, using large language models (LLMs), to automatically answer assurance protocol questions based on evidence extracted from multiple sources.
- **Rule-based Systems and ML Classifiers** to flag non-compliance, identify missing data, and suggest corrective actions.

The AI must be capable of reasoning across **heterogeneous, multimodal inputs**, often with incomplete, inconsistent, or domain-specific terminology. It must then map this information to the assurance scheme protocols—typically a complex, dynamic framework of hundreds of questions and compliance checks.

Key questions for the workshop:

- 1. How can we assure the robustness, consistency, and context-awareness of multimodal AI systems operating across diverse data sources?**
- 2. What evidentiary standards must AI meet to ensure that its answers to assurance questions are auditable, transparent, and aligned with regulatory interpretation? How do these compare with current practices for human inspectors?**
- 3. How do we ensure that AI outputs can be validated, challenged, or corrected by human users—without undermining trust or introducing new risks?**
- 4. How does such a system adapt to new regulations and standards over time, which are sometimes rapidly changing?**

Case Study 3: AI-Assisted Detection of Infections and Other Pre/Post-Mortem Pathologies in UK Abattoirs

As part of an ongoing drive to enhance food safety and operational efficiency, software developers and equipment manufacturers are examining the use of machine learning (ML) to detect signs of infection, and even quality defects, through image recognition. Traditionally, this role is carried out by trained meat inspectors and official veterinarians, who visually assess carcasses for signs of disease or contamination.

AI systems developed by technology providers will deploy high-resolution imaging and ML models trained on thousands of annotated images to identify visual markers of infection and anomalies in real time. These systems would integrate:

- **Deep Learning** to detect visual anomalies or markers of pathology in real-time images or video streams.
- **Transfer Learning** to adapt pre-trained models for new or under-represented pathologies.
- **Edge AI** to enable on-device, real-time decision-making within abattoir environments.

In general, based on the known performance of "deep" learning detection of features in images, these systems are likely to have high potential and could well be scaled in industry at pace. In controlled environments and conditions, they might well demonstrate recall and precision rates comparable to those of human inspectors; where recall (true positives / (true positives + false negatives)) equates to a measure of detection rate and precision (true positives / (true positives + false positives)) to accuracy. However, variability in lighting, carcass presentation, environment, biological diversity, and rare pathogen manifestations are likely to remain key challenges. There are also important questions about the transparency and traceability of AI decisions, particularly in the context of regulatory compliance and public trust.

Key questions for the workshop:

- 1. How can regulators ensure that AI systems for infection detection achieve - and maintain - accuracy, recall, and stability at a level equivalent to or exceeding that of trained human inspectors?**
- 2. What standards and validation processes should be established to evaluate the diversity and quality of training data, especially for rare or emergent pathogens where symptoms may not be well represented in existing datasets?**
- 3. What criteria must be met before AI systems can be authorised for use in regulated environments, and how can ongoing performance be**

assured, particularly in terms of drift, bias, or unforeseen failures?

Case Study 4: AI-Powered Document Inspection at UK Ports of Entry

The UK Food Standards Agency along with software developers are exploring the use of AI systems, including large language models (LLMs), to enhance the inspection and verification of food import documentation at ports. These documents, ranging from health certificates and commercial invoices to packing lists and shipping manifests, are critical for ensuring food safety, regulatory compliance, and traceability of goods entering the UK. Traditionally, official controls involve officers manually reviewing these documents to assess conformity with safety standards and detect inconsistencies or fraudulent entries. This process can be time-consuming, especially under increased trade volumes and complex global supply chains.

An AI-based solution, incorporating document classification models, optical character recognition (OCR), and LLMs, could increase the productivity of frontline officers (e.g., freeing up time for more physical inspections and investigations). These systems can automatically extract key data points, cross-check documents for internal consistency, flag anomalies or incomplete submissions, and even interpret unstructured or multilingual content. LLMs, specifically, have shown promise in identifying subtle discrepancies in language, such as ambiguous product descriptions or suspicious edits.

To address these complexities, AI developers are likely to explore the use of multiple systems to automate document inspection. These systems would integrate:

- **Optical Character Recognition (OCR)** to digitise printed or handwritten documentation.
- **Document Classification Models** to identify and categorise incoming paperwork.
- **Large Language Models (LLMs)** for extracting, interpreting, and cross-validating information from unstructured or multilingual text.

- **Retrieval-Augmented Generation (RAG)** which only generates AI outputs that relate to retrieved external, or policy-specific, sources of truth, including published regulations and policy guidance.
- **Anomaly Detection Algorithms** to flag inconsistencies, duplications, or signs of fraud.

AI-enabled systems may well increase throughput with higher detection rates of noncompliant documentation, but questions remain as to how these systems can be assured, especially given the risk of hallucinations, bias, misinterpretation due to language subtleties, and over-reliance on AI outputs in critical decision-making.

Key questions for workshop:

- 1. How can we ensure the accuracy, reliability, and auditability of LLMs when used to assess official import documentation in regulated environments?**
- 2. What safeguards are required to manage the risk of false positives or negatives, omissions, or AI-generated hallucinations, especially when decisions impact food safety or border clearance?**
- 3. How do we validate training data diversity and alignment with UK regulatory terminology, languages, and document formats to ensure equitable and robust performance?**
- 4. What level of accuracy by an AI system is considered acceptable and would result in reducing burden on human inspectors? Should an AI system result be binary (pass or not) or should it provide a more nuanced output that includes reasons for an assessment?**

AI Glossary

Anomaly Detection Algorithms

Techniques that identify outliers or irregularities in data; flagging errors, fraud, or noncompliance in complex documentation or operational workflows.

Deep Learning

An advanced form of machine learning using layered neural networks to recognise complex patterns; likely effective in tasks such as image recognition within carcass inspection.

Document Question Answering (DocQA)

An AI capability that allows systems to answer specific questions based on the contents of documents, useful for automating certification responses or audit checks.

Edge AI

AI that runs locally on devices rather than in the cloud, enabling real-time decision-making in operational settings like abattoirs or port inspections.

Food Assurance

Processes and schemes that provide verified confidence to consumers, regulators, and businesses that food has been produced, processed, and handled according to defined standards relating to safety, quality, animal welfare, and environmental impact.

Food Safety

Activities and measures aimed at protecting consumers from foodborne illnesses and contamination by ensuring that food is safe to eat. This includes the prevention, detection, and management of biological, chemical, and physical hazards throughout the food supply chain, in line with statutory requirements enforced by the Food Standards Agency.

Food Safety Management Plans

Structured documentation outlining processes, controls, and evaluations to ensure food products are safe for consumption and comply with regulatory standards. These plans incorporate assessments of chemical, microbial, and allergenic risks.

Knowledge Graphs

Structured networks that represent relationships between entities; such as ingredients, allergens, or contaminants, helping AI reason about risk pathways and regulatory linkages.

Large Language Models (LLMs)

AI systems trained on extensive text data that can interpret, summarise, and generate natural language. In food safety and assurance, they could be used to process regulatory documents, inspection records, and policy guidance.

Machine Learning (ML)

A core approach in AI where models learn from data to detect patterns and make predictions or decisions, used widely across safety assessment, document analysis, and image inspection.

Multimodal AI

AI systems capable of processing and integrating multiple data types, such as text, tables, images, and numerical values. This is especially useful in contexts where data is fragmented or presented in different formats.

Natural Language Processing (NLP)

A subfield of AI focused on understanding and interpreting human language potentially used to extract insights from farm records, safety reports, or multilingual documents.

Optical Character Recognition (OCR)

Technology that converts scanned, printed, or handwritten text into machine-readable data, allowing AI systems to work with legacy forms or paper-based documentation.

Retrieval-Augmented Generation (RAG)

A method where AI retrieves relevant external documents before generating a response, ensuring outputs are grounded in verifiable sources such as regulations or guidance notes. For example, if asked, “What are the UK import requirements for soft cheeses?”, a RAG-enabled system would first retrieve current FSA or border import guidelines, then generate a response based specifically on that content, thereby reducing the risk of error or hallucination.

Rule-Based Systems

AI tools that operate using fixed logic rules (e.g., “if X and Y occur, trigger a warning”), providing predictable outputs and supporting regulatory logic or protocol adherence.

Shelf-Life Data

Information generated from studies that determine how long a food product remains safe and

of acceptable quality under defined storage conditions. Critical for labelling, safety assessments, and regulatory compliance.

Third-Party Certification and Assurance

Independent verification processes where external bodies assess farms or food businesses against specific standards related to food safety, environmental protection, and good agricultural and processing practices.

Traceability

The ability to track the history, application, or location of a food product through all stages of production, processing, and distribution. Essential for rapid response to food safety incidents and regulatory compliance.