Science Council Report of Project 'Artificial Intelligence Applications in Food Safety and Authenticity'

Introduction

In this guide

In this guide

- 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

Artificial Intelligence (AI) technologies are rapidly transforming both society and industry, reshaping how food is produced, selected and consumed. While the pace of technological development and adoption presents major commercial opportunities, it raises important questions about the safe and responsible use of AI across the food system. The objectives of this FSA Science Council report are to anticipate likely impacts of known and emerging AI systems and to assess potential implications for food safety and assurance. It considers how AI could affect critical food safety functions and explores perspectives on the standards required for AI function. This report represents the FSA's first formal examination of AI in the food system. However, we acknowledge that the risks and perceptions of this diverse and rapidly evolving technology will become clearer as adoption increases and its intended, and unintended, consequences are known.

Modern AI represents a family of technologies that includes machine learning, computer vision, robotics, natural language processing, and large language model typologies (LLMs), each offering distinct capabilities for tasks such as detecting foodborne risks, automating visual inspections, interpreting regulatory documents, translating multilingual records, and extracting insights from complex or unstructured data. Applications may engage a single AI function or multiple interconnected technologies, for example, combining computer vision for image recognition, machine learning for pattern detection, and large language models (LLMs) for interpreting documentation or generating decision support. The evolution of these technologies has occurred at astonishing pace. Key breakthroughs include AlexNet (Krizhevsky et al. 2012), where deep convolutional neural networks dramatically improved image recognition accuracy over classic computer vision techniques, and large language models (LLMs), highlighted with the launch of ChatGPT as recently as November 2022 (OpenAI, 2022), which brought generative AI into widespread public view and commercial use.

Al adoption across the UK food system is accelerating, particularly in food manufacturing, logistics, and primary production, where technologies such as machine learning, computer vision, and robotics offer productivity gains. In manufacturing, Al is likely to improve production efficiency, safety compliance and quality control; potentially detecting non-conformances more reliably and at greater scale than manual checks, while also reducing labour costs and waste. Newer approaches, such as imitation (Li et al., 2025) and reinforcement learning, will enable robots to mimic complex and dextrous human behaviours found in harvesting, handling, and inspection tasks that are beyond current state-of-theart machines. These advances support not only greater automation but also improved responsiveness to changing supply chain conditions and consumer demands. As Al systems become more accessible and interoperable, they are expected to underpin a shift toward more adaptive, data-driven decision-making across the entire food system.

Al has the potential to transform food safety by shifting assurance processes from largely reactive responses to more proactive, predictive and real-time management. Emerging applications include predictive analytics to anticipate pathogen risks (Benefo et al., 2022), computer vision systems that can continuously monitor processing environments for hygiene and non-conformances (Zhao et al., 2025), and digital twins that model facility operations to optimise preventive controls (Pennells et al., 2025). Al has demonstrated potential to improve both the chemical and microbiological safety of food. Machine learning has shown the potential to improve source attribution in foodborne outbreaks

when combined with whole genome sequencing (Munck et al., 2020). A recent review by Kabir et al. (2025) suggested that machine learning applied to hyperspectral imaging data had potential to classify grains and nuts according to mycotoxin contamination. A significant body of research is already available in this area using many different ML algorithms. Al can also extend surveillance beyond the factory floor, using natural language processing to analyse consumer complaints or social media signals, and machine learning to integrate disparate datasets into early-warning systems for contamination or fraud (Tao et al., 2021). Al could also be used to both generate but also detect fraudulent activity (e.g. fake documents, fake certificates, fake labels etc); positively it could enable regulators and food businesses to respond faster, and reduce reliance on sampling or retrospective testing.

Al offers opportunities to enhance consistency and scale. Unlike human inspectors who must work within time and resource limits, AI systems can continuously scan large volumes of multimodal data (images, text, sensory data etc) across production lines or even supply chains, potentially identifying trends and anomalies invisible to individual auditors. By supporting human decision-making with richer evidence, AI could reduce variability between inspectors, increase sampling rates whilst enabling more transparent traceability from farm to fork. Whilst many consider these applications remain unevenly developed, they highlight the direction of travel: food safety may become more predictive, more integrated, and more responsive as AI tools mature. In short, the technology has the capacity to enhance both the efficiency and resilience of assurance systems, provided it is deployed with the appropriate safeguards. However, the translation of these capabilities into real-world settings must be handled carefully. Depending on how AI is introduced, it could lead to significant changes in worker roles or perceived redundancy of certain tasks, raising serious concerns with jobs and needs for reskilling.

At the same time, the integration of AI into food safety and assurance raises fundamental questions about accountability, explainability, and trust. The FSA operates within a robust legal and regulatory framework, underpinned by the Food Safety Act (1990), which places ultimate responsibility for food safety on the human decisions made by employees and directors of Food Business Operators (FBOs). This accountability cannot be transferred to an algorithm. As the Law Commission (2025) has recently highlighted, autonomous and adaptive AI systems "do not currently have separate legal personality... [and] could lead to 'liability gaps', where no natural or legal person is liable for the harms caused" (Law Commission, 2025). This concern is highly relevant to the food system,

where unexplained or unverifiable AI outputs could undermine both consumer protection and the due diligence defence relied upon by FBOs.

These considerations mean that the deployment of AI in food safety must be accompanied by clear governance, transparency, and human oversight. The challenge for the FSA is to balance innovation and efficiency with regulatory assurance, ensuring that AI augments rather than replaces the human accountability that underpins food law.

Recent research has emphasised that ethical considerations are inseparable from the deployment of AI in food systems. Manning et al. (2022) argue that adoption of AI will only be trusted if it is grounded in a shared vocabulary of ethical principles that stakeholders across the supply chain can understand and apply. Their review identifies seven interlinked aspects, transparency, traceability, explainability, interpretability, accessibility, accountability and responsibility, as central to embedding AI in food governance. Importantly, they highlight that failure to differentiate or operationalise these aspects risks creating barriers to adoption, undermining trust and amplifying bias. For regulators such as the FSA, these findings underline that the introduction of AI in food assurance is not simply a technical question but also a socio-ethical challenge: AI must be explainable, accountable and accessible in ways that align with existing food safety responsibilities if it is to support, rather than erode, consumer confidence (Manning et al., 2022).

Complementing this ethical perspective, Qian et al. (2023) highlight the breadth of Al applications emerging in food safety. They emphasise that adoption remains limited compared to other areas of the agri-food system, constrained by fragmented data sharing, privacy and commercial sensitivity concerns, lack of standardisation, and the absence of clear legal frameworks. Many systems remain at the research stage, often product- or pathogen-specific, with limited scalability into operational practice. Addressing these barriers will require investment by businesses in digital infrastructure, harmonisation of data standards, and frameworks that safeguard both privacy and regulatory integrity.

Taken together, these studies reinforce a common conclusion. Al will be most effective in food safety as a decision-support system operating under human oversight, embedded in strong ethical, legal and governance structures, rather than as a replacement for human accountability. The risks of Al in the food system are not confined to technical performance; they extend to how accountability is assigned, how outputs are explained, and how governance mechanisms maintain oversight.

While AI offers real opportunities to enhance food safety, there is a parallel risk that overstatement or hype could undermine trust in the technology. If inflated claims are allowed to dominate, they risk damaging the reputation of AI before its genuine benefits can be realised. There have been many published cautions about the importance of separating hype and exaggerated claims about AI from the reality (Huckins, 2025 Shoham, 2025). Commentators have shared cases where AI has resulted in unexpected outcomes. To date, most incidents have been relatively small, but they argue for an open but cautious approach. The advent of so called "agentic AI" appears to be at the high end of the hype curve currently and it has been pointed out that, as yet, there is no shared definition of an "agent" in AI (Shoham, 2025). However, AI agents are characterized by combining the power of AI (e.g. LLM) with the specificity of a task (e.g. booking a ticket). In a food safety context, deployment of similar tools would need robust guardrails and close supervision. It is not currently possible to foresee where such a tool could be deployed in food assurance.

Whilst it is becoming clear that AI has transformational power within the food system, its safe and effective adoption will depend on addressing a set of persistent challenges. Questions of data quality, interoperability, transparency, accountability, and legal liability remain central, and the balance between innovation and assurance will be critical. For the FSA, this means considering not only how AI might strengthen food safety controls, but also how its deployment might create new risks, shift responsibilities, or alter the operation of due diligence defences under existing law.

To explore these issues in depth, this report draws on a series of case studies examining the deployment of AI by FBOs across diverse food safety contexts, including product risk assessments, certification and assurance audits, pathology detection in abattoirs, and import documentation checks at ports of entry. Whilst the focus of the report was on FBO application of AI, there is no doubt some of the tools deployed are likely to assist regulators. The case studies, while not exhaustive, provided insights into both the opportunities and risks of AI, illustrating where the technology might augment food safety processes, where it might complicate them, and what governance principles will be needed to ensure safe, fair and trusted adoption across the UK food system.