

Building an Auditable Artificial Intelligence Layer in an Agri-Food SME: A Business Case Study of Universal Organic Products Ltd

Mukesh Kumar Singh¹, Hemant Sharma²

¹ AI Marketplace & Commercial Development, Universal Organic Products Limited, London, United Kingdom

² Market Development & Commercial Operations, Universal Organic Products Limited, London, United Kingdom

ARTICLE INFO

Received:04 Aug 2026
Accepted:22 Sept 2026
Published: 04 Oct 2026

ABSTRACT

Small and medium-sized enterprises (SMEs) have been rapidly embracing artificial intelligence (AI), but most adoption remains limited to fleeting applications and is not well mapped. Most dot-com – can be now artificial intelligence (AI) adoption among small and medium-sized corporations (SMEs) is minimal and shallow, and was even less well-documented. This paper describes a single case study of a UK start-up marketplace for organic products, Universal Organic Products Ltd (UOP), which in 2026 adopted an integrated AI and automation layer and began producing a formal evidence pack for external review. Based on document analysis of technical explanation documents, workflow documentation, and platform screenshots, together with audit-layer counts of platform records, the study analyzed the design, execution, and proof of how eleven functions – ranging from marketplace assistant and allergy-aware recommendations to supplier matching and fraud review – were enabled by the use of AI. The common theme is that rather than any specific skill, UOP's most unique feature is an evidence-first architecture that ensures all AI decisions are saved into specially created audit tables, that data on the original platform is explicitly separated from content generated by the LLM, and that the data is not able to be backdated. This design serves three purposes: regulatory traceability, investor scrutiny, and organizational learning. The paper provides a grounded description of AI adoption in response to disparate evidentiary pressures, suggests evidence-first AI adoption as a dynamic capability for SMEs, and highlights risks associated with toxic products such as special-category health data, as well as allergen liability and cold-start marketplace risks.

Keywords: Artificial intelligence Adoption; small and medium-sized enterprises; agri-food e-commerce; multi-sided platforms; AI auditability; AI governance; case study

1. INTRODUCTION

The food supermarket sector's strongest part is certified organic food, with UK sales reaching around £3.9 billion in 2025, marking 14 years of continuous growth. 83% of households buy organic, but only 3% of land is farmed organically, relying heavily on imports (Soil Association Certification, 2026). Global retail sales were €145.0 billion in 2024 (FiBL & IFOAM – Organics International, 2026). India is a key exporter, worth USD 665.97 million in 2024-25 (Press Information Bureau, 2025). The UK-India Free Trade Agreement, effective from July 15, 2026, eliminates tariffs on Indian organic foods, boosting demand among qualified farmers aiming to reach UK buyers. Meanwhile, AI adoption in UK businesses has grown from 12% in 2023 to about 35% in June 2026, mainly through large language models. 50% of firms are not using AI extensively, down from 76% in 2023 (Office for National Statistics, 2026). SMEs face greater challenges than larger firms in finance, skills, data, compliance, and funding access. There is a lack of detailed studies on AI implementation in micro-enterprises, with current research limited to self-reported data and no detailed info on AI architecture or processes. Accountability is also an issue, with increasing regulation requiring AI disclosures and record-keeping (European Parliament & EU Council, 2024), but specific guidance on compliance for SMEs remains scarce. This paper examines the shortfalls through a case study

of a multi-sided marketplace-build-as-a-service provider undergoing organic certification. The startup, known as a "stand-up or fall" marketplace, faced evidentiary pressure in May 2026, prompting founders to build a formal evidence pack with AI and automation, claimed to be 'actively deployed.' The study addresses three questions: how the AI layer was designed and implemented, influenced by technology, organizational, and environmental factors; how its evidence-first design serves as a strategic capability beyond compliance; and how the AI layer impacts innovation, viability, scalability, and risks. The paper shares an artifact-level account of AI adoption, introduces evidence-first AI, maps AI features to regulatory responsibilities, and brainstorms residual risks.

2. LITERATURE REVIEW AND ANALYTICAL FRAMEWORK

2.1 Organic Markets, Platforms and Credence Goods

Organic food's defining traits, like certification, can't be verified through inspection or tasting, relying instead on third-party certification (Darby & Karni, 1973). Soil Association is the main certifier in Britain, and imported products need a certificate from the exporter's control body, certified in Britain, plus an importer and consignee (Department for Environment, Food & Rural Affairs, 2024). Certification and import controls establish institutional hierarchy by raising entry costs but also provide trust signals that add value. The UK organic market nearly doubled in 10 years, with a 4.2% growth in 2025 (Soil Association Certification, 2026). Riverford Organic Farmers, an established, employee-owned grower and box scheme, is a leading direct-to-consumer platform, projecting about £117 million in 2025 (Bell & Barthold, 2024; Companies House, 2025). Conversely, Farmdrop, a low-margin marketplace backed by venture investors, collapsed in December 2021 due to funding issues and failure to attract paying customers, leaving producers unpaid (Grocery Gazette, 2021). The key strategic difference is between a successful integrated grower and an asset-focused marketplace, which UOP must address. Multi-sided platforms enable direct customer interactions, with cross-side network effects: more users increase value (Hagiu & Wright, 2015; Parker & Van Alstyne, 2005; Rochet & Tirole, 2006). They face a chicken-and-egg challenge and must build trust, offering mechanisms for validation, listings, and dispute resolution. Many ventures start as resellers before becoming platforms, highlighting this progression..

2.2 AI Adoption in SMEs and Agri-Food E-Commerce

The literature on technology adoption mainly focuses on Tornatzky and Fleischer's (1990) TOE framework, which states that adoption depends on technology, organizational factors like size and management support, and the external environment, including competitors and regulations (Baker, 2011). This is often paired with Rogers' (2003) Diffusion of Innovations theory, where perceived attributes (relative advantage, compatibility, complexity, trialability, observability) influence diffusion speed. Sánchez et al. (2025) combine these to study AI adoption by SMEs, addressing issues like data access and governance, and adapting the framework for responsible AI. Ayinaddis (2025), reviewing 78 studies, finds SMEs less prepared technologically, financially, and for compliance, noting that the TOE framework can obscure relationships among these areas. Another perspective sees AI as a strategic asset. The resource-based view highlights valuable, rare, inimitable, and non-substitutable resources (Barney, 1991), while the dynamic-capabilities view emphasizes routines like sensing, seizing opportunities, and reconfiguring assets. Mikalef and Gupta (2021) describe AI capability as orchestrating tangible resources, human skills, and intangible assets like coordination and change capacity. The intangible aspects are most accessible for micro-enterprises, such as culture, discipline, and design, which are easier to afford than data centers. Generative AI exacerbates one risk. The technology will be fluent but sometimes delusional, as surveys increasingly agree; this is a hallmark of LLMs (Huang et al., 2025; Ji et al., 2023), and users often believe AI-provided content even when it is incorrect (Klingbeil et al., 2024). For SMEs, that means using trusted internal data as the basis for their outputs, reserving meaningful human interaction for high-impact decisions, and falling back to an acceptable level of functionality if the external model breaks down. In agri-food, Reitano et al. (2025) trace AI use along the chain, warning most benefits are perception-based or modeled correlations, not measurable results. Samoggia et al. (2021) find most digital tools in consumer, farmer info, and supply chain networking. The main use case is consumer-oriented: health-aware recommendations for multi-objective systems optimizing preference, health, explanations, and diversity (Zhang et al., 2025). An 'allergy-aware' need raises reliability issues, and studies emphasize presenting these as decision support, not medical advice (Iskrenova-Ekiert & Musa, 2026; Tran et al.,

2021). Machine-learning fraud detection is effective, but reviews point out challenges like class imbalance, concept drift, interpretability, and human oversight (Mutemi & Bacao, 2024).

2.3 Responsible AI, Auditability and Regulation


The regulatory expectations have been solidifying quickly. The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) lays down obligations for logging and record keeping on high-risk systems (Arts. 23-24) and also mandates that deployers ensure that the person with whom they are dealing is aware they are communicating with an AI system (Art. 50). The governance, logging and human oversight component (Art. 12 and 18) requires deployers to keep logs and human oversight for the AI, the need for staff AI literacy (Art. 4), and SMEs to be prioritized in sandbox access (Art. 62) (European Parliament & Council of the European Union, 2024). While this will only affect a UK company when its output is used by customers within the Union, the Act is effectively a reference architecture for accountable AI. The UK has embarked on a principles-based, regulator-led approach (Department for Science, Innovation and Technology, 2023), followed by a growth-oriented action plan since January 2025 (Department for Science, Innovation and Technology, 2025). However, health information such as allergies, food sensitivity or intolerance due to medical conditions, medical notes, etc. is special category data under the UK GDPR as implemented under the Data Protection Act 2018 (DPA) and will require an impact assessment if the processing is likely to be high risk (Information Commissioner's Office, 2023).

There is one more layer: under the Food Information Regulations 2014, 14 regulated allergens must be listed, and in the case of a distance sale, all allergen information must be made accessible before a sale and at the time of delivery. (Food Standards Agency 2026; R474/001: Food Information Regulations 2014). A provision that provides allergy advice, then, centers on a protective accountability matrix around information responsibilities. Between these escapes, the necessity of having a food marketplace that is auditable, meaning that at the end of the day, the system may be shown what it had observed, done, and then decided, is a condition of acceptable gameplay.

2.4 Analytical Framework

The study is organized around TOE, and Rogers' (2003) technological context framework explains the perceived attributes of the designs. The second research question is examined from resource-based and dynamic capabilities perspectives (Barney, 1991; Mikalef & Gupta, 2021). The third question is addressed directly in the company's investor-oriented business plan and reflects what the company expects to be measured against: innovation, viability, and scalability (Universal Organic Products Ltd, n.d.). Figure 1 summarises the framework.

Figure 1 Analytical Framework of the Study



Note. TOE contexts shape the form that AI adoption takes; the resulting architecture is interpreted through resource-based and dynamic-capability lenses and assessed against the innovation, viability, and scalability criteria used in the company's investor-facing business plan.

3. METHOD

3.1 Design and Data

The study uses an instrumental case design (Stake, 1995), although it is not unique in that it examines UOP itself; rather, it highlights how a micro-enterprise adopts AI under external evidentiary pressure. For a single case, it is a good option when it is revelatory or sufficiently rich to trace possible mechanisms in detail (Siggelkow, 2007; Yin, 2018). In the latter case, complete internal documentation for implementing AI in an SME is rarely available at a level of granularity suitable for this kind of study.

Primary data are documents. Document analysis is a systematic task that extracts meaning by reading documents and considering both their content and the context in which they were created (following Bowen 2009). The corpus includes two technical explanation documents prepared on 20 May 2026; five walkthrough documents (timestamped screenshots); a Key Investor Data Document (investor document); and a 66-slide document presenting the evidence with commentary. These were established in conjunction with the Companies House register and sources of current market and policy data, replacing the 2018-2019 statistics from the company's own business plan.

Table 1 Documentary Sources Analyzed

Source	Type	Date	Use in analysis
AI and Automation Features Explanation	Internal technical and evidence document	20 May 2026	Feature inventory; implementation timeline; accuracy rules; evidence-strength grading
Technical Explanation of AI and Automation Features	Developer-side technical document	20 May 2026	Mechanisms; database records; page-level location of features; one-order evidence bundle
Workflow documentation	Five screenshot-annotated walkthroughs of recorded platform use	May 2026	Functional verification of journeys; identification of fallback versus completed model calls; self-reported improvement lists
Project evidence presentation commentary	Slide-by-slide commentary on a 66-slide evidence deck	2026	Cross-reference of storefront, administrative, database, hosting and version-control evidence
Key Investor Data Document	Investor-facing business plan	Undated (pre-platform)	Business model; market positioning; employment plan; statements on innovation, viability and scalability
Platform screenshots	Screen captures from recorded walkthroughs and the evidence deck	May 2026	Visual evidence of implemented functions (Figures 4 to 9)
AI evidence runner acceptance counts.	Administrative console display	May 2026	Aggregate counts of audit records by function (Table 3)
Companies House register	Public record	Accessed 2026	Incorporation date; status; SIC code; registered office
Market and policy publications	Published reports and official guidance	2024–2026	Contextual triangulation; replacement of outdated statistics in company documents

3.2 Analysis, Ethics and Reflexivity

There were three passes of analysis. To create the structure and source for each document, I first did a skim. A close reading then pulled out each AI or automation function, how it worked, and the resulting outputs and evidence tables, leading to the inventory shown in Table 2 and a chain of evidence from each sign back to a document. Interfering between data and theory, in the manner of Eisenhardt (2009), a third, interpretive pass coded the data against the TOE contexts and Rogers' attributes, and then, following the Gioia methodology (Gioia et al., 2013), from the first order of concepts in the company's own terms – 'evidence runner', 'fallback automation', 'accuracy note' – to the second order themes, of which evidence-first adoption was most significant. The analysis rule applied throughout was that a feature had to be implemented if the documents stated the mechanism, location within the platform, and evidence records, and the feature had to be considered an intention if it appeared aspirationally.

Including the case company's name was intentional, as anonymization would have hindered readers' ability to verify public facts and would have reduced practitioner value. Naming also creates responsibilities, and the company agreed to be named in its own documentation; additionally, we never name individuals, and we verified them against readily available public records whenever possible. Two points should be made with caution about the corpus. They are advocacy documents for external reviewers and not performance measures – they exist and how they function – rather than how well they will perform in a live trading environment. As a result, the analysis assumes they constitute valid evidence of design and implementation and that they are claims, not evidence of commercial outcomes. In particular, the documents take the same stance, consistently distinguishing what is proved from what is assumed.

4. THE CASE: UNIVERSAL ORGANIC PRODUCTS LTD

4.1 Company Background and Business Model

Universal Organic Products Ltd, incorporated in England (company number 13566556), is registered in Slough, Berkshire. It's active, with last CIS received in May 2026 (Companies House, 2026). Its SIC code is 56290 (other food services), which does not reflect its digital nature. It has a founder-led management team; one director has over ten years' experience on certified organic farms in India, covering cultivation, regulatory affairs, crop management, and inventory control.

The business aims to create a multi-sided marketplace accessible via website and apps (Universal Organic Products Ltd, n.d.). It connects consumers, retailers, and wholesalers seeking certified organic produce with farmers, producers, and delivery agents. It rejects single-brand box schemes and site-only retailers like Riverford and Eversfield, viewing them as aggregators with multiple suppliers offering live availability. Suppliers must provide organic certification (per the Soil Association) and avoid using non-approved inputs. It plans to import organic products directly from Indian farmers, including oils, seeds, cereals, pulses, spices, tea, coffee, dry fruits, and ayurvedic items. Growth is planned from three founders to about eleven staff.

By 2026, the platform had become a website using PHP and MySQL with PayPal, featuring a storefront, search, customer registration (including health terms), partner dashboards, an admin back office, and a mobile-responsive design (Universal Organic Products Ltd, 2026c). Development included version control, hosting, and backups.

4.2 The AI and Automation Layer

The company recorded, and if required gathered under one umbrella, eleven AI and automation functions between 17 and 20 May 2026 (Universal Organic Products Ltd, 2026a, 2026b). Table 2 outlines the processes, main outputs and evidence logging for each function. These fit into 4 families of customer-facing intelligence (assistant, recommendations, health & allergy guidance); supply-side intelligence (supplier matching); transactional risk and operations (fraud review, plus delivery, billing, marketing & inquiry automation); and infrastructure - the knowledge base with embeddings and the audit dashboard with its evidence runner.

Table 2 AI and Automation Functions Documented by UOP, May 2026

Function	Documented mechanism	Principal output	Evidence records
Smart Marketplace Assistant (LLM chatbot)	Retrieves product, supplier, cart, order, policy, and support context; sends a controlled prompt to the AI service or to a fallback response generator	Structured answer with source evidence and safety note	ai_interactions; ai_chat_conversations; ai_support_messages
Product recommendations	Scores products on profile signals, category fit, organic verification, allergy and diet signals, embeddings, and model reranking where configured	Ranked recommendations with stated reason	ai_recommendation_events; ai_customer_insights
Health and allergy guidance	Compares customer allergies, dietary preferences and medical notes with product allergy flags, dietary tags, keywords and categories	Suitability label; allergy caution; safety guidance	ai_health_events
Supplier matching	Scores suppliers by certification, service area, delivery capability, business type, verified listings, stock, and category or product fit	Selected, backup, and rejected suppliers with explanation and confidence score	ai_supplier_match_events; vendor_order_assignments; procurement_requests
Fraud and payment risk review	Reviews order amount, item count, quantities, contact details, guest status, payment capture status, and fulfillment mode	Risk score and level (low, medium, high); explanation; recommended review action	ai_fraud_events; ai_admin_review_queue; order_risk_fields
Delivery automation	Generates delivery summaries and review items for supplier-routed orders	Delivery summary; review queue item	ai_automation_jobs; ai_admin_review_queue; delivery_order_assignments
Billing automation	Supports checkout and PayPal capture; records billing notes and payment status	Payment status; billing automation note	ai_automation_jobs; order payment fields
Targeted marketing	Uses cart and order signals to propose marketing suggestions	Marketing suggestion	ai_automation_jobs
Contact and inquiry automation	Classifies inquiry intent and priority; drafts summary and next action	Summary; recommended next action	ai_contact_intakes; ai_automation_jobs
Knowledge base and embeddings	Indexes products, suppliers, categories, delivery, billing, and fraud rules, policies, and health guidance as knowledge chunks; generates embedding vectors when an embedding API is configured	Retrieval context for the assistant and recommendations	ai_knowledge_chunks; ai_embeddings
AI audit dashboard and evidence runner	Displays API status, model information, feature counts, and recent records; runs controlled developer test journeys; exports CSV summaries	Audit view; evidence exports	All ai_tables


The design features are recurring. All AI responses are based on platform actions: The assistant's prompt includes product, supplier, cart, order, policy, and support contexts, while recommendations and health functions reference

customer profile information and catalog allergen and dietary tags. If no external OpenAI-compatible model is set up or available, the platform creates an audited, rule-based fallback response and labels it accordingly. Each function logs a structured record with feature, input context, output, source, and timestamp in a size-optimized table, which can be cross-checked with a recommendation screenshot, the relevant row, and the resulting order..

4.3 The Audit and Evidence Architecture

This architecture (Figure 2) involves customer and partner touchpoints feeding into a context-retrieval layer; a controlled prompt is sent to the AI service, which responds with either an external model or an audited fallback, then outputs are sent back to the platform. Decisions are managed in an audit layer with an admin dashboard showing API status, model info, feature counts, recent data, and an evidence runner for controlled AI journeys and CSV exports. Medium- and high-risk fraud assessments and review items in delivery or procurement go to an admin review queue for a man-in-the-middle decision before fulfillment.

Figure 2 UOP's AI Service and Audit Architecture as Documented in May 2026



Note. Compiled from Universal Organic Products Ltd (2026a, 2026b). Numbered blocks follow the order in which requests are processed. Table names are those used in the company's database.


Three accuracy rules govern the architecture: fallback automation is not to be presented as an external LLM output statement until recorded in the log; AI cannot facilitate backdated records; and evidence must relate to actual platform use, not demonstration data (Universal Organic Products Ltd, 2026a, 2026b). Some documents grade evidence strength: the chatbot as strong if only interaction status shows resolved and API billing evidence; recommendations if screenshots or event rows are shown; supplier matching if records of supplier, product, and order are shown. A recommended set of ten artifacts should follow each order, covering the entire customer transaction—from identifying the customer and recommendations, to store access, cart details, transaction routing, payment capture, order receipt, dashboard views, and database changes.

4.4 Documented Workflows and Platform Evidence

Five workflow documents present screen-recorded walkthroughs with screenshots, each with a snapshot of the screen at each stage of the process to which they correspond, including customer and partner registration, the full customer order flow through to administrative order visibility, the assistant, and a wholesaler journey (Universal

Organic Products Ltd, 2026c). As illustrated in Figure 3, the customer order flow is reconstructed, and the points where the AI service is called and audit records are written are identified. The ten steps include four that call the AI service: recommendations, cart suitability, checkout health advice, and fraud review, with fraud review potentially passing the order to a human for final review before payment.


Figure 3 Customer Order Flow With AI Touchpoints and Audit Logging



Note. Reconstructed from the customer order-flow and evidence documentation (Universal Organic Products Ltd, 2026b, 2026c). Shaded steps marked AI invoke the AI service; blue lines indicate records written to the audit layer. Medium- and high-risk orders are held in the administrative review queue before payment proceeds.


Figures 4 to 7 illustrate evidence used by customers. The storefront pumps the assistant as a non-stop widget, preceded by a Live context label (4a), followed by a fixed set of actual products as the outcome of a context-driven call (4b), which the assistant returns based on the shopper's profile information. peringkat's conversation identifier and its mode (external AI), along with the labels of the live sources used (5 products, 1 supplier, 6 categories, 10 Knowledge chunks). This disclosure is at the interface level in accordance with the level 2 grounding and leveling rules. The health & allergy profile page (Figure 5a) lets the customer indicate allergens, medical conditions, and dietary preferences, and the resulting recommendations include an AI match score, the recommendation source, and a plain-language explanation that references the allergens and preferences (Figure 5b).

Figure 4 Storefront and Smart Marketplace Assistant



Note. (a) Public storefront (universalorganic.uk) with the Smart Marketplace Assistant widget open and its live-context label. (b) Assistant conversation in which a profile-based product request returns a summary grounded in five catalog products, with the conversation identifier, mode (external AI), and the live sources used printed beneath the answer—screenshots supplied by the company, May 2026.


Figure 5 Customer Health Profile and AI Product Recommendations



Note. (a) Health and allergy profile page where customers record allergens, medical conditions, and dietary preferences, with the health-conscious shopper summary and recommended products generated from the saved profile. (b) Recommended products showing the AI match score, the recommendation source, and an explanation referencing the customer's dietary preferences and allergens. Screenshots supplied by the company, May 2026.

AI recipes show ingredient suitability in the cart, with a watch note for safety (Figure 6b). They include a profile summary, safety note, and suitability levels, highlighted in the cart (Figure 6a). At checkout, this guidance is added to the basket and billing form, emphasizing it is not medical advice (Figure 6b). After order creation, a risk-review panel appears above payment options (Figure 7), rating it low risk and detailing an order pattern, a health-conscious cart, and valid contact info for automatic approval; medium- and high-risk assessments generate a review queue.

Figure 6: AI Suitability Check in the Cart and Health Guidance at Checkout



Note. (a) Cart page with the AI diet, health, and recommendation check summarising the shopper's profile, a watch note and a safety note, and a per-item suitability label. (b) Checkout page with billing details and AI health guidance attached to the basket, including the statement that the guidance is not medical advice—screenshots supplied by the company, May 2026.


These documents have two analytically important characteristics. They include the names of a displayed response and fallback logic, in accordance with the accuracy rules. They list recommended improvements, such as grouping sections and adding validation in the partner form, among others, that would make the application more accurate, showing that the documents have been used for internal quality review and as external evidence.

4.5 Administrative Evidence and Records in the Audit Layer

The audit layer, accessible via the administrative console, features a scorecard for AI products with an AI audit page listing recommendation evidence. It includes hashes of customer identifiers and queries or contexts, ensuring an

evidence trail (Figure 8a). Health and allergy evidence list customer and order identifiers, preserving anonymity for data minimization (Figure 8b). An evidence runner page (also in Figure 8b) enables the team to create labeled test records without backdating or identifying paid orders, exporting summaries as CSV files and comparing record counts to internal targets.

Figure 8 Administrative AI Audit Dashboard and Evidence Runner



Note. (a) AI audit page listing recommendation evidence (hashed customer identifier and query) and health and allergy evidence (customer, order reference and product identifiers). (b) Evidence runner page showing acceptance counts by evidence area against the company's internal targets; the page states that developer test records are clearly labeled, are not backdated, and do not mark any order as paid. Screenshots supplied by the company, May 2026; browser and operating-system elements removed.

This evidence from the runner shows acceptance counts in Table 3 for May 2026. Three points follow. First, 362 of 376 assistant interactions involved external model calls, with 13 falling back to deterministic automation. External models handled 96% of calls, with few returning to automation. Second, distribution is uneven: recommendation events (312) and health events (26) are more common than other events like supplier matches, fraud, and marketing suggestions (4 each). Health and recommendation events occur on each page view with an associated profile, while other events occur once per order. Third, the embedding count—less than the 50+ target—indicates the knowledge-base embedding pipeline was used but not widely, as reflected in the counts from logs. These show which functions are used and how often, but not customer demand.

Table 3 Records Logged in the Audit Layer as Displayed by the Evidence Runner, May 2026

Evidence area	Records logged	Internal acceptance target	Corresponding audit record
Completed interactions	362	10 or more	ai_interactions (status: completed)
Fallback interactions	13	Monitor	ai_interactions (source: fallback automation)
Embeddings	4	50 or more	ai_embeddings

Recommendations	312	3 or more	ai_recommendation_events
Supplier matches	4	3 or more	ai_supplier_match_events
Health events	26	3 or more	ai_health_events
Fraud events	4	2 or more	ai_fraud_events
Automation jobs	25	8 or more	ai_automation_jobs
Marketing suggestions	4	2 or more	ai_automation_jobs (marketing)
Review queue items	13	4 or more	ai_admin_review_queue

Note. Counts as displayed on the platform's administrative evidence runner page (Figure 8b). They include clearly labeled developer and test records as well as records from ordinary platform use, and therefore evidence that the logging functions operate and how intensively each function was exercised during the documentation period, not customer demand. Targets are the company's internal acceptance thresholds.

The administrative product register (Figure 9) shows that 217 products have been matched to a supplier, with details of category, verification of the product as organic, price, pack size, stock level, and date added. When routing business-to-business, bulk, and checkout orders (Table 2), the engine must be able to score these fields to work with real data rather than illustrative data.

Figure 9 Supplier and Product Verification Records Used by the Supplier-Matching Engine



Sno	Image	Name	Supplier	Category	Description	Organic	Price	Size	Stock	Added At	
1		Organic Banana	Fresh Organics	Fresh Fruits & Vegetables	Fresh organic bananas supplied for everyday retail and wholesale fruit demand. Carefully selected for reliable ripeni...	Verified	£27.15	Case	120	14 Mar 2026, 04:03 AM	
2		Organic Apple	Fresh Organics	Fresh Fruits & Vegetables	Crisp organic apples chosen for balanced sweetness, clean bite and dependable shelf display. Suitable for snacking, l...	Verified	£24.20	Case	110	14 Mar 2026, 04:03 AM	
3		Organic Tomato	Fresh Organics	Fresh Fruits & Vegetables	Vibrant organic tomatoes with a juicy interior and clean natural acidity. Ideal for salads, sandwiches, fresh cooking...	Verified	£11.70	1kg	90	14 Mar 2026, 04:03 AM	
4		Organic Oranges	Fresh Organics	Fresh Fruits & Vegetables	Sweet and fragrant organic oranges packed for fresh retail movement, juicing and everyday household use. Bright colour...	Verified	£14.68	Case	100	14 Mar 2026, 04:03 AM	
5		Organic Onions	Fresh Organics	Fresh Fruits & Vegetables	Kitchen staple organic onions with firm texture and full savoury flavour. Suitable for roasting...	Verified	£7.16	1kg	140	14 Mar 2026, 04:03 AM	

Note. Administrative product register showing 217 products matched to a supplier, with category, organic verification status, price, pack size, stock level, and date added. These fields feed the supplier-matching scores described in Table 2—screenshot supplied by the company, May 2026.

5. FINDINGS

5.1 TOE Analysis

The design incorporates four of Rogers' (2003) attributes, with high compatibility due to the AI being embedded in functions like the storefront, shopping cart, checkout, and admin pages—rather than as a stand-alone app (Figures 5-7). A retrieval-then-prompt paradigm ensures the model stays linked to platform data, following hallucination literature (Huang et al., 2025). Trialability is built through the evidence runner, which tests functions in controlled

journeys, preventing untested behavior; observability uses the audit dashboard to show real-time customer activity to management (Figure 8; Table 3). The hybrid architecture external model, when configured properly, emphasizes cost, reliability, and credibility. External API calls are charged per use, and failures require fallback to deterministic recommendations to avoid errors. The firm's AI promises credibility, making rule-based outputs unsafe; thus the external model acts as a trusted fallback, with 362 of 375 interactions documented for model calls (Table 3). Much platform intelligence relies on rule scoring, acknowledged in their evidence grading documentation. UOP, with limited resources, benefits from strong funding and domain expertise in organic cultivation and India supply networks. Their disciplined approach to evidence, documentation, and accuracy, reinforced by an embedded evidence runner, creates a cost-effective, accountable AI resource. However, Slack remains a significant constraint., with some features dependent on API setups and the expected usage path not being achieved.

Three environmental forces influenced adoption. The first is market opportunity: a fourteen-year growth market for organics in the UK; one that relies on imports; and one that will now be served by a trade agreement to lower tariffs on the same categories the company wishes to provide products for in India (House of Commons Library, 2025; Soil Association Certification, 2026). The second is regulation: organic import regulations, the allergen law, health-data laws, and rules for transparency regarding AI all affect the platform, each with its own reasons for keeping records. The third and most unique is external evidentiary scrutiny; by creating an evidence pack for reviewers to consider in relation to innovation, viability and scalability, the AI documentation has been rolled up. In TOE terms, this is the outside pressure that did not just influence acceptance, but the type of adoption: being required to prove what had been developed led to an architecture that could be proven. Competitive pressure rounds out the picture, too: competitors' positioning relative to integrated box schemes will drive differentiation, and Farmdrop's demise is the failure mode they are trying to avoid. Table 4 summarises.

Table 4 Summary of TOE Analysis

Context	Factors identified	Illustrative evidence
Technological	High compatibility with existing platform; complexity managed through retrieval-then-prompt grounding; trialability via evidence runner; observability via audit dashboard; hybrid external-model and fallback architecture	AI panels embedded in cart, checkout and administrative pages; controlled prompt built from platform context; developer test journeys; labeled fallback records
Organisational	Founder-led micro-enterprise; deep organic-supply domain knowledge; limited development and data-science capacity; strong documentation and accuracy discipline; human-in-the-loop review	Supplier scoring on certification and service area; accuracy notes; evidence-strength grading; prohibition on backdating; administrative review queue
Environmental	Growing, import-reliant UK organic market; UK–India trade liberalization; organic import, allergen, data-protection and AI transparency regimes; external evidence review; integrated incumbents and a failed marketplace precedent	Market data 2024–2026; GB certificate-of-inspection regime; investor-facing criteria of innovation, viability and scalability; positioning against box schemes; Farmdrop collapse

5.2 Evidence-First AI Adoption as a Dynamic Capability

The second research question is whether evidence architecture is a strategic capability or a compliance cost. The dynamic capabilities view suggests it is both, and that it is valuable for three functions. It helps to sense – the dashboard and evidence runner provide information about which features are being used, which suppliers are matched and rejected, how often the fraud review ‘trips’ and where it fails – something a micro-enterprise does not have the resources or workforce to be able to gather. It facilitates “seizing” of new features which can be exercised and proven before release – decreased cost of experimentation. It also enables switching between fallback and

model paths: Labeling and logging of the fallback and model paths allow the firm to easily switch between the two and adjust either switching providers or altering the scoring rules while maintaining continuity of evidence.

The structure also means that each operation creates data. Every suggestion event, supplier vetting, and fraud detection contributes to a systematic record of buyer preferences, supplier reliability, and risk trends unique to each platform, and it builds over time. In a multi-sided market, this information can benefit the network in proportion to its growth, but so far it is scarce, and its merits remain to be determined. The conclusion is that UOP has created, from the asset's inception, the mechanism by which it will be formed, and one which is beneficial to multiple stakeholder groups: regulators, investors, and managers, which satisfies several at the same time: traceability, proof, and feedback. We call this evidence-first AI adoption because the logic and requirements to log, label, and review human work become part of the design process rather than a post-mortem report of what the AI system has done.

5.3 Assessment Against Innovation, Viability and Scalability

Table 5 compares the venture against the three criteria, focusing on evidence in the corpus and on gaps. In the context of innovation, the three features (the certified-supply marketplace, personalization with allergies, and auditability on the AI layer) are unique in a sector with single-brand, integrated models. None of the features are particularly new, though, so the innovation lies in the integration and accountability, not necessarily the algorithm. Purely on viability, the platform is built, workflows shown end to end – with payments even live captured – and expertise on the supply side is real. However, the corpus has no amount of trading volume, unit economics, or retention. On the model, the system is capital-light compared with an inventory-holding grower; its audit system scales with volume; the India corridor is now more conducive, and scaling requires addressing the chicken-and-egg dilemma on both sides, plus per-consignment compliance costs for cross-border organic imports.

Table 5 Assessment Against Innovation, Viability and Scalability Criteria

Criterion	Supporting evidence in the corpus	Gaps and risks
Innovation	Integration of a certified-supply marketplace, health- and allergy-aware personalization, supplier matching and an auditable AI layer; distinct from integrated single-brand incumbents; India sourcing corridor	Individual components have precedents in the literature; novelty rests on integration and accountability architecture rather than on algorithms; mobile applications described in the plan are evidenced mainly through responsive web layouts.
Viability	Functioning platform with end-to-end demonstrated journeys including live payment capture; founder domain expertise in organic cultivation and Indian supply; documented development, hosting and backup effort	No evidence in the corpus of trading volume, unit economics, or retention; business-plan market data predate the platform; evidence strength for several AI functions depends on API configuration
Scalability	Capital-light marketplace model relative to inventory-holding growers; audit architecture designed to scale with volume; employment plan to eleven roles; favorable trade environment after the UK–India agreement	Chicken-and-egg problem on both sides of the platform; per-consignment organic import compliance cost; margin arithmetic on low-value baskets; employment trajectory not yet realized

Note. Criteria as framed in the company's investor-facing business plan (Universal Organic Products Ltd, n.d.).

6. DISCUSSION

6.1 Theoretical and Managerial Implications

The case adds to TOE literature by showing how environmental pressures influence not just the likelihood of adopting new technologies but also the form. Existing research treats regulatory/institutional aspects as barriers or enablers (Ayinaddis, 2025; Sánchez et al., 2025). Evidence from external reviews affected architecture because logging, labeling, and reviewing were necessary; UOP would otherwise have deferred it. This supports Sánchez et al. (2025) that responsible AI governance should be integrated, with governance supported by investment, review components, or regulators. The case also shows replacing tangible resources with AI as an intangible resource. Mikalef and Gupta (2021) state AI capability relies on material, intangible, and human resources. UOP is strong in material, medium in intangible, and weak in human resources; its accuracy model improves trust and adaptability, aligning with dynamic capabilities theory. The case exemplifies the shallow yet broad AI adoption in the UK (Office for National Statistics 2026), demonstrating that deep AI use is less about tech sophistication and more about integrating it into daily practices and management. It's not just about providing a transferable pattern but giving that pattern through the case itself. AI answers should build on the company's data, explained in the interface, with fallback options. Record all decisions and outcomes for review, and route significant ones for human judgment. Clearly articulate accuracy rules, evaluate evidence, and build comprehensive testing. These steps generate documentation preferred by investors, auditors, and regulators. The evidence bundle offers a template with interconnected artifacts to show how AI systems work.

6.2 Regulatory and Ethical Risks

Table 6 shows how each AI function connects to the regulatory considerations it raises and the documented design response. The following 3 risks are emphasized.

The first concerns health data. The health-profile feature, as a special category of data, falls under the GDPR in the United Kingdom, which is defined as 'allergies' or 'medically related dietary restrictions' or 'free text medical notes', which must have a clear lawful basis, most likely 'explicit consent', with limits on minimization, retention, and, due to the sensitivity, limits on processing, which must also include a Data Protection Act 2018 Data Protection Impact Assessment (Council of Europe, 2023). They constitute safety guidance, not medical advice, and they are recorded in an events table designed for this purpose, which is appropriate. However, the output resulted from no completed impact assessment and no consent mechanism – both of which are priority actions for the corpus.

The second is about the allergen liability. Before purchase and at delivery, the platform (not just its suppliers) must provide allergen information under the two-stage information duty in the Food Information Regulations 2014. An AI suitability label, based on information encoded by the advertiser/supplier, cannot be relied upon if the label is incorrect and has had a reassuring effect on the customer or customers, and then the conditions which gave the guidance its legal meaning would be misleading. Mitigations include keeping a log of the tag values used for each health event, checking the underlying declaration displayed with any label for AI, or keeping the fourteen regulated allergens as a structured field on the supplier listing.

The third is on transparency and oversight. The assistant should introduce itself as an AI system at the start of every interaction, as required by Article 50 of the EU AI Act of the European Parliament & Council of the European Union (2024). All staff executing fraud scores and supplier matches must also have AI literacy, as the Act mandates for all staff involved in its deployment. Finally, the review queue should include not only a machine recommendation, but also a record of the human decision made, as a machine recommendation alone does not demonstrate oversight. The architecture is well suited to these aspirations, and the latter, which is more of procedure than of architecture, remains at the base of the Royal Palace.

6.3 Market Risk: Cold Start and the Farmdrop Lesson

There are no significant technical threats. Network-effects theory forecasts that a marketplace with few producers will see few consumers, and vice versa (Parker & Van Alstyne, 2005; Rochet & Tirole, 2006). It cannot create demand, but it is a supply-side "bootstrap"--UOP's supplier-matching engine will route an order to a suitable

certified supplier even if that supplier's catalog is limited. Platform sequencing research indicates that many platforms that succeed as resellers for a curated range open up after which to third parties (Hagiu & Wright, 2015), which can be the approach made as the India import channel if costs of the per-consignment certificate-of-inspection and risks of extra controls applied to product-origin combinations are priced in (Department for Environment, Food & Rural Affairs, 2024).

The lesson of Farmdrop's collapse is clear: insurance- and logistics-tech-fueled business models with very low grocery margins, and new-business offers that customers adopted as a supplement to their shopping regime but not a replacement for it (Grocery Gazette, 2021). UOP's design is intentionally lighter because a supplier/delivery partner places and fulfills orders, but margin arithmetic means any platform charging a commission on low-value baskets is subject to the same rules. For the first time, the evidence architecture will enable the company to see basket value, fulfillment costs, and repeat rate as they happen. That will be judged on those numbers, and not on how sophisticated the artificial intelligence might be.

7. LIMITATIONS AND FUTURE RESEARCH

The study has four main drawbacks: it relies on documents, not interviews or observation, so it doesn't verify real-time behavior during design and implementation. These documents were created for external review and tend to be advocacy, but their accuracy mitigates this. The corpus lacks effectiveness measures, offering only aggregate record counts, preventing claims about AI effectiveness beyond existence, design, and operation. A single case supports analytical, not statistical, generalization (Yin, 2018). These limitations suggest future research: extending the architecture for longitudinal follow-up using export data to measure feature usage, fallback rates, recommendation uptake, and fraud detection over time. User testing of guidance on health and allergy topics can address framing, understanding, and explainability issues noted in the literature (Iskrenova-Ekiert & Musa, 2026; Tran et al., 2021). Comparing case studies of SMEs under different external scrutiny will evaluate how scrutiny influences AI use. Finally, a log-labeling tool that records evidence, labels, reviews by humans, self-grades, and tests could correlate results across firms.

8. CONCLUSION

Universal Organic Products Ltd aims to connect certified organic producers in Britain and India with UK trade buyers and consumers. By May 2026, it plans to deploy an AI layer impacting every customer and supplier interaction, including a context-aware assistant, health suggestions, recommendations, supplier matching, fraud review, and operational automation. Its architecture links AI decisions to platform data, records context and outcomes, and involves human oversight where necessary, designed to satisfy regulators, funders, and managers. The case shows external review influences SME AI adoption and accuracy discipline, balancing tangible resources with AI. It offers a template for accountable AI for small businesses, guiding questions for evidence and intelligence development. UOP's goal is to prioritize AI profitability, with future data revealing if this approach is effective.

REFERENCES

- [1] Ayinaddis, S. G. (2025). Artificial intelligence adoption dynamics and knowledge in SMEs and large firms: A systematic review and bibliometric analysis. *Journal of Innovation & Knowledge*, 10(3), 100682. <https://doi.org/10.1016/j.jik.2025.100682>
- [2] Baker, J. (2011). The technology–organization–environment framework. *Information Systems Theory: Explaining and Predicting Our Digital Society*, Vol. 1, 231-245.. https://doi.org/10.1007/978-1-4419-6108-2_12
- [3] Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- [4] Bell, E., & Barthold, C. (2024). Managing employee ownership transitions for sustainability in food and farming enterprises: Learning from Riverford: the Open University and Sustain.
- [5] Bowen, G. A. (2009). Document analysis as a qualitative research method. *Qualitative Research Journal*, 9(2), 27-40. <https://doi.org/10.3316/QRJ0902027>

- [6] Companies House. (2025). Riverford Organic Farmers Limited (Company number 03731570): Filing history and accounts. <https://find-and-update.company-information.service.gov.uk/company/03731570>
- [7] Companies House. (2026). Universal Organic Products Ltd (Company number 13566556): Company overview. <https://find-and-update.company-information.service.gov.uk/company/13566556>
- [8] Darby, M. R., & Karni, E. (1973). Free competition and the optimal amount of fraud. *The Journal of law and economics*, 16(1), 67-88.. <https://doi.org/10.1086/466756>
- [9] Data Protection Act 2018, c. 12 (UK). <https://www.legislation.gov.uk/ukpga/2018/12>
- [10] Department for Environment, Food & Rural Affairs. (2024). Importing and exporting organic food. GOV.UK. <https://www.gov.uk/guidance/importing-and-exporting-organic-food>
- [11] Department for Science, Innovation and Technology. (2023). A pro-innovation approach to AI regulation (CP 815). GOV.UK. <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>
- [12] Department for Science, Innovation and Technology. (2025). AI Opportunities Action Plan. GOV.UK. <https://www.gov.uk/government/publications/ai-opportunities-action-plan>
- [13] Eisenhardt, K. M. (2009). Building theories from case study research. *Studi organizzativi*, (2008/2)..http://euroac.ffri.hr/wp-content/uploads/2010/06/Eisenhardt_1989_Building-Theories-from-Case.pdf
- [14] European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024/1689. <http://data.europa.eu/eli/reg/2024/1689/oj>
- [15] FiBL & IFOAM – Organics International. (2026). The world of organic agriculture: Statistics and emerging trends 2026 (27th ed.). Research Institute of Organic Agriculture FiBL and IFOAM – Organics International.
- [16] Food Information Regulations 2014, SI 2014/1855 (UK). <https://www.legislation.gov.uk/ukxi/2014/1855>
- [17] Food Standards Agency. (2026). Allergen guidance for food businesses. <https://www.food.gov.uk/business-guidance/allergen-guidance-for-food-businesses>
- [18] Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational research methods*, 16(1), 15-31.. <https://doi.org/10.1177/1094428112452151>
- [19] Grocery Gazette. (2021, December 22). Suppliers speak out about Farmdrop collapse. <https://www.grocerygazette.co.uk/2021/12/22/suppliers-speak-out-about-farmdrop-collapse/>
- [20] Hagiu, A., & Wright, J. (2015). Multi-sided platforms. *International journal of industrial organization*, 43, 162-174. <https://doi.org/10.1016/j.ijindorg.2015.03.003>
- [21] House of Commons Library. (2025). UK–India Free Trade Agreement (Research Briefing CBP-10258). UK Parliament. <https://commonslibrary.parliament.uk/research-briefings/cbp-10258/>
- [22] Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., ... & Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM transactions on information systems*, 43(2), 1-55..<https://doi.org/10.1145/3703155>
- [23] Information Commissioner's Office. (2023). Guidance on AI and data protection. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>
- [24] Iskrenova-Ekiert, E. K., & Musa, M. (2026, July). User-Centered Design for Allergy-Aware Food Recommendation System. In *International Conference on Human-Computer Interaction* (pp. 438-451). Cham: Springer Nature Switzerland.. https://link.springer.com/chapter/10.1007/978-3-032-31418-5_27
- [25] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., ... & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12), 1-38.. <https://dl.acm.org/doi/pdf/10.1145/3571730>

- [26] Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and reliance on AI: An experimental study on the extent and costs of overreliance on AI. *Computers in Human Behavior*, 160, 108352. <https://www.sciencedirect.com/science/article/pii/S0747563224002206>
- [27] Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, 58(3), 103434. <https://www.sciencedirect.com/science/article/pii/S0378720621000082>
- [28] Mutemi, A., & Bacao, F. (2024). E-commerce fraud detection based on machine learning techniques: Systematic literature review. *Big Data Mining and Analytics*, 7(2), 419-444. <https://ieeexplore.ieee.org/iel7/8254253/10506765/10506811.pdf>
- [29] Office for National Statistics. (2026). Artificial intelligence in UK businesses: 2023 to 2026. ONS.
- [30] Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management Science*, 51(10), 1494-1504. <https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=1177443>
- [31] Press Information Bureau. (2025, December 18). Export of organic food products [Press release]. Ministry of Commerce and Industry, Government of India.
- [32] Reitano, M., Segovia, M. S., & Nayga Jr, R. M. (2025). A systematic review on the impact of artificial intelligence in the agri-food supply chain. *Food Policy*, 137, 102983. <https://www.sciencedirect.com/science/article/pii/S0306919225001885>
- [33] Rochet, J. C., & Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3), 645-667. https://publications.ut-capitole.fr/id/eprint/1184/1/rochet_tirole.pdf
- [34] Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- [35] Samoggia, A., Monticone, F., & Bertazzoli, A. (2021). Innovative digital technologies for purchasing and consumption in urban and regional agro-food systems: A systematic review. *Foods*, 10(2), 208. <https://www.mdpi.com/2304-8158/10/2/208>
- [36] Sánchez, E., Calderón, R., & Herrera, F. (2025). Artificial intelligence adoption in SMEs: Survey based on TOE-DOI framework, primary methodology and challenges. *Applied Sciences*, 15(12), 6465. <https://www.mdpi.com/2076-3417/15/12/6465>
- [37] Siggelkow, N. (2007). Persuasion with case studies. *Academy of Management Journal*, 50(1), 20-24. https://faculty.wharton.upenn.edu/wp-content/uploads/2012/05/AMJ_Persuasion.pdf
- [38] Soil Association Certification. (2025). Organic market report 2025. Soil Association.
- [39] Soil Association Certification. (2026). Organic market report 2026. Soil Association.
- [40] Stake, R. E. (1995). *The art of case study research*. SAGE.
- [41] Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington Books.
- [42] Tran, T. N. T., Felfernig, A., Trattner, C., & Holzinger, A. (2021). Recommender systems in the healthcare domain: state-of-the-art and research issues. *Journal of Intelligent Information Systems*, 57(1), 171-201. <https://link.springer.com/content/pdf/10.1007/s10844-020-00633-6.pdf>
- [43] Universal Organic Products Ltd. (n.d.). Key investor data document [Unpublished investor-facing business plan].
- [44] Universal Organic Products Ltd. (2026a). AI and automation features explanation [Unpublished internal document, 20 May 2026].
- [45] Universal Organic Products Ltd. (2026b). Technical explanation of AI and automation features [Unpublished internal document, 20 May 2026].
- [46] Universal Organic Products Ltd. (2026c). Platform workflow documentation: Customer order flow, user registration, partner registration, marketplace assistant and wholesaler journey [Unpublished internal documents].

- [47] Yin, R. K. (2018). Case study research and applications (Vol. 6, pp. 200-224). Thousand Oaks, CA: sage..
- [48] Zhang, Z., Wang, Z., Ma, T., Taneja, V. S., Nelson, S., Le, N. H. L., Murugesan, K., Ju, M., Chawla, N. V., Zhang, C., & Ye, Y. (2025). MOPI-HFRS: A multi-objective personalized health-aware food recommendation system with LLM-enhanced interpretation. In Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25) (pp. 2860–2871). ACM. <https://doi.org/10.1145/3711896.3737420>