

Optimised AI Limited

Human Oversight Mechanisms

Documentation of how and when humans review AI outputs, particularly for high-risk applications.

Document Owner	Optimised AI Limited — Board of Directors
Trading Name	GetFrugal (www.getfrugal.uk)
Version	1.0
Effective Date	March 2026
Review Cycle	Annual (or upon material change to AI systems)

1. Purpose

This document describes the human oversight mechanisms in place at Optimised AI Limited (trading as GetFrugal) to ensure responsible deployment and monitoring of our AI-powered spending analysis platform. It details when, how, and by whom AI outputs are reviewed, and the escalation procedures in place for edge cases.

2. Risk Classification

GetFrugal's AI system is classified as a low-risk application under the framework of the EU AI Act and UK Government guidance on AI regulation. Our rationale is as follows:

- The AI analyses publicly available deal data and user-provided transaction records to generate money-saving recommendations
- It does not make autonomous financial decisions, execute transactions, or move money on behalf of users
- It does not provide regulated financial advice; all outputs are informational recommendations only
- Users retain full control over whether to act on any recommendation
- No decisions with legal or similarly significant effects are made by the AI

Despite the low-risk classification, we maintain robust human oversight as a matter of best practice and to uphold user trust.

3. AI System Architecture and Scope

Our proprietary, in-house AI system performs the following functions:

- **Classifying bank statement entries into spending categories (e.g., groceries, subscriptions, utilities)**Transaction categorisation:
- **Identifying recurring payments that may represent active or forgotten subscriptions**Subscription detection:
- **Matching spending patterns against available deals, cashback offers, and loyalty rewards**Savings opportunity identification:
- **Generating tailored suggestions for reducing expenditure based on the user's actual spending behaviour**Personalised recommendations:

The AI does not: access users' bank accounts directly; initiate payments or cancel subscriptions on behalf of users; provide regulated financial, legal, or tax advice; or make automated decisions that produce legal effects concerning users.

4. Human Oversight Framework

4.1 Pre-Deployment Review

Before any AI model update or new feature is released to production:

- All model changes undergo internal testing against a curated dataset of sample transactions
- Output quality is manually reviewed by a member of the development team to confirm accuracy of categorisations and recommendations
- Edge cases and failure modes are documented and addressed prior to release
- A sign-off is required from both a technical lead and a business stakeholder before deployment

4.2 Ongoing Monitoring

Post-deployment, the following ongoing oversight mechanisms are in place:

- **System monitoring flags anomalies in AI output patterns (e.g., unusually high numbers of flagged subscriptions, miscategorisation spikes, or recommendation accuracy drops)** Automated alerts:
- **A random sample of AI-generated analyses is reviewed by a team member each week to verify quality and accuracy** Weekly output sampling:
- **Users can flag incorrect categorisations or irrelevant recommendations directly within the platform. All flagged items are reviewed by a human within 48 hours** User feedback loop:
- **Key metrics (categorisation accuracy, recommendation relevance, user satisfaction scores) are tracked on an ongoing basis and reviewed in monthly operational meetings** Performance dashboards:

4.3 Escalation Procedures

The following escalation framework applies:

Trigger	Action	Responsible
User reports incorrect recommendation	Manual review within 48 hours; correction applied if confirmed	Customer Support / Technical Lead
Automated monitoring detects anomaly	AI outputs paused for affected function; manual investigation initiated	Technical Lead
Systemic accuracy drop detected	Full model review; rollback to previous version if necessary; incident report filed	Board / Technical Lead

4.4 Human-in-the-Loop for Sensitive Outputs

While the vast majority of GetFrugal's AI outputs are low-risk informational recommendations, certain categories receive additional human scrutiny:

- Recommendations involving potential financial claims or rewards above a specified threshold are flagged for human review before being presented to the user

- Any AI output that could be interpreted as financial advice is reviewed to ensure appropriate disclaimers are displayed
- New recommendation categories or deal types are approved by a human before being added to the system

5. User Control and Autonomy

Users maintain full autonomy over their interaction with GetFrugal's AI:

- All AI-generated recommendations are presented as suggestions, never as instructions
- Users can dismiss, ignore, or provide feedback on any recommendation
- Users can opt out of AI analysis at any time via my.getfrugal.uk/opt-out
- No automated actions are taken on behalf of the user without their explicit initiation

6. Staff Training

All team members involved in the development, deployment, or monitoring of AI systems receive training on: responsible AI principles and bias awareness; the company's human oversight procedures; data protection obligations under UK GDPR; and escalation and incident response protocols. Training is refreshed annually or when significant changes are made to the AI system.

7. Record Keeping

We maintain records of all human oversight activities, including: pre-deployment review sign-offs and test results; weekly sampling logs and findings; user-reported issues and resolution outcomes; escalation incidents and corrective actions taken; and model version history and change logs. These records are retained for a minimum of two years and are available for audit upon request.