

Anti-Slavery Intelligence (ASI): An AI-Powered Tool for Modern Slavery Compliance Analysis and Remediation

Mahmoud Gad¹, Abdessalam Elhabbash¹ and Steven Young¹

¹Lancaster University Management School, UK
{m.gad1, a.elhabbash, s.young}@lancaster.ac.uk

Abstract

We present Anti-Slavery Intelligence (ASI), a deployed AI system that analyses corporate modern slavery statements to identify compliance gaps and generate prioritised remediation advice. ASI orchestrates a multi-model pipeline: Gemini-2.5-Flash for vision-enhanced PDF parsing and structured compliance scoring against 48 expert-defined criteria, and Gemini-2.5-Pro for synthesising company-specific, timeline-based recommendations. A benchmarking engine compares each statement against industry and FTSE index averages across 11 sectors. Evaluation on 95 expert-annotated statements yields an F1-score of 0.86 and recall of 0.94, reducing time to generate an initial compliance assessment from several hours to approximately five minutes per document. ASI is publicly available at <https://www.antislaveryintelligence.co.uk/> with free access for academics and NGOs. A demonstration video is at <https://youtu.be/DNhdeGrWzg>.

1 Introduction

Modern slavery affects over 50 million people globally [ILO *et al.*, 2022]. The UK Modern Slavery Act 2015 requires organisations with turnover above £36 million to publish annual statements describing the steps taken to identify, mitigate, and remediate forced labour and human trafficking risks in their operations and supply chains [UK Government, 2015]. In practice, high-quality statements should explain governance accountability, supply-chain risk assessment, due-diligence processes, training, effectiveness indicators, and responses when cases are found. Studies show that many disclosures remain incomplete, boilerplate, or disconnected from risk management practices [BHRRC, 2016; Nersessian and Pachamano, 2022; Mai *et al.*, 2023]. Modern slavery is not only a disclosure failure: forced labour and trafficking risks arise in real supply chains, especially where outsourced, temporary, or migrant work obscures accountability and workers face coercive practices such as recruitment-fee debt or passport retention [ILO *et al.*, 2022; CCLA, 2024]. These harms make statement quality consequential, because weak disclosures can hide whether firms

are finding cases, remedying affected workers, and preventing recurrence [CCLA, 2024]. Expert review of these statements typically requires several hours per document and is difficult to scale across the thousands of organisations required to report.

Manual assessment creates three practical problems: (i) high time and expertise costs limit the number of statements that can be reviewed; (ii) subjective interpretation leads to inconsistent scoring across reviewers; and (iii) limited cross-industry benchmarking disadvantages smaller organisations that lack specialist advisory services.

We introduce *Anti-Slavery Intelligence (ASI)*, a deployed AI system that addresses these problems through an end-to-end automated pipeline. Given a PDF statement, ASI: (1) extracts content using vision-enhanced layout analysis; (2) scores the statement against a 48-criterion framework derived from CCLA’s “Find It, Fix It, Prevent It” guidelines [CCLA, 2024]; (3) benchmarks performance against sectoral and FTSE index averages; and (4) generates prioritised recommendations structured by implementation timeline (immediate, short-term, and long-term actions).

Existing AI-assisted approaches to modern slavery statement analysis focus on numeric scoring. Poon *et al.* [2025] demonstrate LLM-based classification on a small dataset, while AIMSCheck [Bora *et al.*, 2025] introduces chain-of-thought scoring with cross-jurisdiction transfer. TISCreport [TISCreport, 2024] provides an AI-assisted compliance audit service but, to our knowledge, lacks peer-reviewed evaluation. To our knowledge, none of these systems combine interactive tooling, cross-sector benchmarking, and timeline-based remediation advice in a single publicly accessible platform.

ASI contributes: (a) a *multi-model AI pipeline* that separates parsing, scoring, benchmarking, and feedback generation into specialised stages; (b) a *dynamic benchmarking engine* comparing performance across 11 sectors and 3 FTSE indices; (c) *timeline-based remediation recommendations* that translate compliance gaps into concrete actions; and (d) a *deployed, publicly accessible web application* serving corporate, investor, and NGO stakeholders.

2 System Architecture

Figure 1 shows the end-to-end pipeline, which orchestrates four stages using two LLMs accessed via managed cloud




Figure 1: Multi-model architecture of ASI. Gemini-2.5-Flash handles PDF parsing (Step 1) and compliance scoring (Step 2). A benchmarking engine computes sectoral comparisons (Step 3). Gemini-2.5-Pro generates prioritised remediation recommendations (Step 4).

APIs without fine-tuning.

Step 1: Vision-enhanced parsing. Gemini-2.5-Flash processes the uploaded PDF using its multimodal (vision + text) capabilities, handling complex layouts, embedded tables, charts, and graphics that challenge traditional text extraction.

Step 2: Structured compliance scoring. The same model evaluates extracted content against our 48-criterion framework. Criteria are grouped into batches of 16 per API call, balancing context richness against latency. Internally, each criterion is represented as targeted evidence checks with deterministic aggregation rules, rather than a single free-form scoring prompt. Near-deterministic decoding (temperature 0.05) and content-addressable caching, keyed by a SHA-256 document fingerprint and model version, provides deterministic scores for previously analysed statements and reduces variance for new inputs. Each criterion receives a structured justification and supporting evidence extracted from the source text.

Step 3: Contextual benchmarking. A benchmark engine aggregates statistics from previously analysed statements across 11 industry sectors and 3 FTSE indices (100, 250, 350), computing percentile rankings that show where each company leads or falls behind its peers.

Step 4: Feedback synthesis. Gemini-2.5-Pro translates scores into plain-language narratives and prioritised recommendations grouped into immediate (0–3 months), short-term (3–6 months), and long-term (6–12+ months) actions. This larger model generates company-specific remediation advice contextualised by sector risks, rather than generic suggestions.

Evaluation frameworks. ASI ships with two assessment frameworks. The *Current Framework* (48 criteria) operationalises CCLA’s “Find It, Fix It, Prevent It” methodology [CCLA, 2024], a standard developed with investors managing over £15 trillion in assets. Its checks cover statutory statement formalities, disclosure quality, risk identification, remediation, and prevention. Examples include board approval and senior sign-off, identification of migrant or temporary labour as a risk factor, disclosure of recognised forced-labour indicators such as passport retention, remedy outcomes where issues are found, and adoption of the employer-pays recruitment principle. The *Advanced Framework* (61 criteria) extends this with Transparency in Supply Chains (TISC) and International Reporting Template criteria for organisations seeking deeper compliance maturity assess-




Figure 2: ASI results dashboard for a real FTSE 100 company (company name redacted for anonymity). (a) Overall compliance score (86%) with sector and index median comparisons. (b) Sector benchmark showing 75th-percentile ranking against Consumer Staples peers. (c) Prioritised strategic actions grouped by implementation timeline.

ment. Users select their preferred framework at upload time. Our evaluation (Section 4) uses the 48-criterion framework, which has been validated against expert annotations.

Deployment. ASI is deployed as a cloud-native application with a React single-page application frontend, a serverless Python backend on Google Cloud Functions, and EU-only data residency for GDPR compliance. Passwordless email-link authentication provides secure access, and PDFs are processed in memory with only hashes and structured results persisted. Per-document analysis takes approximately five minutes, though runtime may vary with API latency. The modular pipeline design is model-agnostic: swapping in a newer LLM requires no architectural changes.

3 Demonstration Scenario

The demonstration will walk attendees through ASI’s workflow, with emphasis on usability for non-technical compliance professionals and a self-contained explanation of the technical pipeline behind each dashboard result.

Live demo flow. An attendee selects a modern slavery statement from a FTSE 100 company and uploads the PDF through the web interface, optionally specifying the industry sector, stock index, and assessment framework (48-criterion CCLA standard or 61-criterion advanced framework including TISC international criteria). As the pipeline runs, a real-time progress panel, powered by server-sent events (SSE), displays the current stage (parsing, scoring, benchmarking, feedback synthesis). For previously analysed statements, cached results load in seconds; for new uploads, end-to-end processing typically takes approximately five minutes. The

	Acc.	F1	Prec.	Rec.
Overall	0.82	0.86	0.80	0.94
Core Req.	0.96	0.98	0.98	0.98
Disclosure	0.90	0.94	0.90	0.99
Due Diligence	0.80	0.73	0.64	0.84
Prevention	0.74	0.78	0.69	0.89

Table 1: Performance on 95 expert-annotated statements (binary per-criterion classification). Benchmarking pool: FTSE 350 companies across 11 sectors. Runtime: ~ 5 min/doc.

interactive results dashboard (Figure 2) shows:

- *Compliance scores.* Overall and category-level scores (Core Requirements, Disclosure, Due Diligence, Prevention) displayed as cards and bar charts, with expandable per-criterion rationales and evidence excerpts from the source document.
- *Benchmarking visualisations.* Comparative plots showing the company’s performance against its sector average and FTSE index peers, highlighting areas where the company lags behind or leads.
- *Recommendation roadmap.* Timeline-based remediation advice categorised by theme (e.g., risk assessment, supplier engagement, governance), with each recommendation linked to specific criteria and evidence.
- *Export.* Structured reports in PDF, JSON, or CSV format for internal stakeholders or external advisors.

We will analyse both a high-quality and a low-quality statement to demonstrate how ASI distinguishes substantive disclosures from boilerplate text, and we will compare the ASI output against expert annotations side-by-side.

Interactive elements. Attendees can upload their own PDF statements (or choose from provided examples), select industry sectors, and explore the full results dashboard. The demonstration highlights how different industries exhibit distinct compliance profiles and how the benchmarking feature contextualises performance relative to peers.

4 Evaluation

We evaluate ASI’s implementation of the CCLA 48-criterion framework (45 criteria assessable via document analysis)¹ on 95 real-world modern slavery statements scored by two independent CCLA compliance experts per statement, with disagreements resolved through moderation. Each criterion is treated as a binary classification (met or not met) and metrics are computed across all criterion–statement pairs.

Table 1 summarises performance. The high overall recall (0.94) is critical in compliance settings where missing a material gap is more costly than flagging a non-critical issue. Core Requirements achieve near-perfect scores as they capture binary legal obligations. Due Diligence and Prevention

are more challenging due to complex, narrative descriptions of processes and governance, but recall remains high (0.84–0.89), aligning with our goal of minimising missed gaps.

Ethical Statement

ASI operates in a sensitive human-rights domain. We position the system as a *decision-support tool*: it surfaces likely compliance gaps and suggests improvements, but final judgments remain with human experts and regulators. Automated scoring may introduce biases against organisations with less polished disclosures; the high recall setting mitigates missed gaps at the cost of potential false positives, which require careful human interpretation. Per-criterion justifications and source-text evidence are generated at scoring time and synthesised into the dashboard to support auditability. ASI processes only publicly available statements and does not store personal data; EU-only data residency ensures GDPR compliance.

Acknowledgements

We thank CCLA Investment Management for access to their expert-annotated dataset, and Ewan Warburton and Ranbir Kochar for research assistance. This work was supported by Lancaster University Management School, EPSRC, ESRC, and AHRC.

References

- [ILO *et al.*, 2022] ILO, Walk Free, and IOM. Global Estimates of Modern Slavery: Forced Labour and Forced Marriage. International Labour Organization, Geneva. <https://doi.org/10.54394/CHUI5986>, 2022.
- [UK Government, 2015] UK Government. Modern Slavery Act 2015, Section 54: Transparency in Supply Chains. <https://www.legislation.gov.uk/ukpga/2015/30/section/54>, 2015.
- [CCLA, 2024] CCLA Investment Management. Find it, Fix it, Prevent it: Modern Slavery Report 2024. <https://www.ccla.co.uk/sites/default/files/2024-10/CCLA-Find-it-Fix-it-Prevent-it-Annual-Report-2024-Final-1.pdf>, 2024.
- [Nersessian and Pachamano, 2022] David Nersessian and Dessislava Pachamano. Human trafficking in the global supply chain: Using machine learning to understand corporate disclosures under the UK Modern Slavery Act. *Harvard Human Rights Journal*, 35:515–535, 2022.
- [Mai *et al.*, 2023] Nam Mai, Petros Vourvachis, and Suzana Grubnic. The impact of the UK’s Modern Slavery Act (2015) on the disclosure of FTSE 100 companies. *The British Accounting Review*, 55(3):101115, 2023. <https://doi.org/10.1016/j.bar.2022.101115>.
- [BHRRC, 2016] Business & Human Rights Resource Centre. FTSE 100 At the Starting Line: An Analysis of Company Statements Under the UK Modern Slavery Act. BHRRC, London, 2016.

¹Three criteria require external verification: homepage link to the statement, modern-slavery registry registration, and precise fiscal-year coverage.

- [Poon *et al.*, 2025] Ser-Huang Poon, Eghbal Rahimikia, Philip Jobi Vallavanthra, and Siliang Wei. Analysing modern slavery statements using large language models. In *Handbook of Artificial Intelligence in Higher Education*, chapter 9, pages 134–151. Edward Elgar Publishing, 2025.
- [Bora *et al.*, 2025] Adriana Eufrosina Bora, Akshatha Arodi, Duoyi Zhang, Jordan Bannister, Mirko Bronzi, Arsene Fansi Tchango, Md Abul Bashar, Richi Nayak, and Kerrie Mengersen. AIMSCheck: Leveraging LLMs for AI-assisted review of modern slavery statements across jurisdictions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9109–9135, 2025. <https://aclanthology.org/2025.acl-long.446/>.
- [TISCreport, 2024] TISCreport. Modern Slavery Act Statement AI Audit. <https://tiscreport.org/products/service/modern-slavery-act-statement-ai-audit>, 2024.