Document reference: PO40.00

Version: 2.0

Effective date: 31 October 2026

Next review due: October 2027

1.0 SUMMARY

Magnum GSI Ltd is committed to the responsible, transparent, and ethical use of artificial intelligence (AI) in the delivery of its services. This policy governs how AI technologies are used across Magnum GSI Ltd operations, including the GeoScreen automated geotechnical desk study platform, to ensure AI is deployed safely, lawfully, and in a manner consistent with our professional obligations.

AI systems used by Magnum GSI Ltd do not make autonomous decisions. All AI outputs are subject to human interpretation and professional judgement prior to use. AI is used solely as a decision-support tool to assist qualified professionals and clients in understanding geotechnical risk; it does not replace professional engineering judgement.

This policy is aligned with ISO 42001:2023 (Artificial Intelligence Management System) and ISO 27701:2019 (Privacy Information Management System). It should be read alongside the Data Security Policy (PO9.00), GDPR Policy (PO19.00), and Confidentiality Policy (PO7.00).

2.0 DEFINITIONS

AI System — A machine-based system that processes inputs to generate outputs such as predictions, recommendations, decisions, or content, with varying levels of autonomy.

GeoScreen — Magnum GSI Ltd's proprietary automated geotechnical desk study platform, which uses AI to generate structured professional reports from public geoscientific data sources.

Large Language Model (LLM) — An AI model trained on large datasets capable of generating human-like text. Used within GeoScreen via the Anthropic Claude API.

AI Output — Content generated by an AI system, such as report narrative, summaries, the reading of historical maps, and the conceptual site model. In GeoScreen, hazard ratings are not AI outputs: they are set by explicit, documented rules applied to the source data. Outputs may be incomplete or omit relevant site-specific constraints due to limitations in underlying datasets and model interpretation, and should be treated as indicative only.

Human Review — Assessment of AI outputs by a qualified person prior to professional, engineering, or planning reliance. Human review must be proportionate to risk and complexity. Where outputs inform

engineering, planning, or financial decisions, review by a suitably qualified geotechnical professional shall be required.

Data Processor (AI) — A third-party organisation that processes personal or project data on behalf of Magnum GSI Ltd as part of an AI pipeline (e.g. Anthropic PBC).

Prompt Governance — The rules, constraints, and instructions applied to AI model inputs to ensure consistent, accurate, and professionally appropriate outputs.

AI Risk Register — A controlled document recording identified risks associated with AI systems in use, including likelihood, impact, and mitigation measures. Reviewed at least annually and following any material system change, with risks scored based on likelihood and impact. Maintained by Technical Management and approved by the Director.

3.0 SCOPE

This policy applies to all AI systems used by Magnum GSI Ltd, all employees, directors, freelancers, and subconsultants working on behalf of Magnum GSI Ltd, and all third-party AI service providers engaged by Magnum GSI Ltd.

The primary AI system currently in scope is the GeoScreen platform and its integration with the Anthropic Claude API for automated report generation. Any new AI systems must be assessed against this policy and recorded in the AI Systems Register before deployment.

AI outputs may be incomplete or omit relevant site-specific constraints due to limitations in underlying datasets and model interpretation. All outputs should be treated as indicative only and are not a substitute for a professionally authored report or intrusive ground investigation where site conditions require it. This limitation is disclosed to clients in the Limitations section of every GeoScreen report (Section 12 of Tier 1, Section 8 of Tier 0).

4.0 AI SYSTEMS IN USE

The following AI Systems Register records all AI systems in active use. It is maintained as a controlled document, reviewed at least annually and following any material system change, with risks scored based on likelihood and impact.

AI System / Integration	Purpose and Data Processed	Human Oversight
GeoScreen — Anthropic Claude API (claude-sonnet-5-5 from October 2026, or successor models).	Generation of report narrative and conclusions (Tier 1), the Tier 0 key takeaway and the conceptual site model; reading of historical maps for features on and around	Report reviewed by client or qualified geotechnical engineer prior to engineering or planning reliance.

Provider: Anthropic PBC. Data transfer: US infrastructure; lawful under UK Addendum to EU SCCs.	the site (two independent readings must agree before an on-site feature affects a rating); and the automated QA check of each report before issue. Hazard ratings themselves are set by explicit rules, not by AI. Data processed: site address/postcode and public geoscientific API outputs. No special-category personal data submitted.	Limitations section included in every report. Client notified of AI use in the report's Limitations section (Section 12 of Tier 1, Section 8 of Tier 0).

5.0 OBJECTIVES

In alignment with ISO 42001:2023, Magnum GSI Ltd commits to the following objectives for responsible AI governance:

- Deploy AI systems only for defined, legitimate, and documented purposes as recorded in the AI Systems Register.
- Ensure AI systems do not make autonomous decisions; all AI outputs are subject to human interpretation and professional judgement prior to use.
- Ensure all AI outputs that inform professional, engineering, or planning decisions are subject to appropriate human oversight before reliance.
- Maintain transparency with clients about the role of AI in service delivery, including disclosure in all AI-generated reports, presented in plain language suitable for non-technical users.
- Conduct and document AI risk assessments for each AI system in use, reviewing them at least annually.
- Protect personal and project data processed through AI systems in accordance with UK GDPR and ISO 27701:2019.
- Monitor AI outputs for systematic bias across different site types, geological conditions, and geographic areas.
- Ensure no client data is used to train third-party AI models, enforced through contractual controls with AI service providers.
- Maintain and review prompt governance rules to ensure AI outputs remain accurate, consistent, and professionally appropriate.
- Achieve continual improvement in AI governance through monitoring, internal audit, and annual management review.

6.0 DATA GOVERNANCE AND PRIVACY (ISO 27701)

The use of AI systems at Magnum GSI Ltd is governed by the following privacy principles, consistent with UK GDPR and ISO 27701:2019:

Data Minimisation: Only the minimum data necessary to generate the requested AI output is submitted to AI processors. GeoScreen submits the site location (address, postcode, grid reference or drawn outline), the proposed use and the public data retrieved for the site; customer email addresses and payment details are never sent to the AI provider; no special-category personal data is intentionally processed within the AI pipeline.

Purpose Limitation: Data submitted to AI systems is used solely for the generation of the requested report or output and for no other purpose.

Processor Controls: AI service providers are engaged under appropriate Data Processing Agreements (DPAs). Anthropic's DPA incorporates the UK Addendum to the EU Standard Contractual Clauses, providing a lawful transfer mechanism under UK GDPR for data processed on US infrastructure.

No Training on Client Data: Magnum GSI Ltd ensures, via contractual controls, that client or project data submitted to AI APIs is not used to train or fine-tune AI models.

Automated Decision-Making: GeoScreen produces AI-assisted outputs that inform professional judgement. These outputs do not constitute solely automated decision-making with legal or similarly significant effects on individuals within the meaning of Article 22 of UK GDPR. Human review is required prior to any consequential reliance.

Data Residency: AI processing via the Anthropic Claude API currently uses global (US-routed) infrastructure. Transfer is lawful under the UK Addendum to EU SCCs. EU inference routing will be requested where commercially available post-launch.

7.0 AI RISK MANAGEMENT (ISO 42001)

Magnum GSI Ltd maintains an AI Risk Register as a controlled document. Risk assessments are reviewed at least annually and following any material change to an AI system in use, with risks scored on a likelihood-impact matrix. The register is maintained by Technical Management and approved by the Director.

Accuracy and Reliability: Regular quality assurance testing of AI outputs against a validated gold set of reference reports. Outputs are verified against explicit prompt governance rules before release.

Automated QA Gate: Every report is checked before issue by rule-based checks and an AI review against the language and content rules. A report that fails is held and sent to Magnum for review instead of to the customer.

Hallucination and Factual Error: AI outputs are generated from structured, verified data inputs (public API sources) and governed by explicit prompt rules that restrict unsupported statements. All reports include a limitations section advising professional review.

Bias and Fairness: AI model outputs are monitored for systematic bias across different site types, geological conditions, and geographic areas. Anomalies are investigated and prompt governance updated.

Security of AI Credentials: API keys and credentials are stored in encrypted environment variables, rotated regularly, and are never committed to version control. Access is restricted to authorised systems.

Service Availability: Fallback procedures are in place for AI service interruptions, including pipeline error handling and client notification of delays.

Scope Creep: AI is used only for the purposes defined in the AI Systems Register. Any material change to system purpose, input data type, or output type shall trigger a formal review under this policy and an update to the AI Systems Register, with prior written approval from the Director required before deployment.

Third-Party Model Changes: Material changes to AI model behaviour by providers (including model version updates) are monitored through QA testing. Prompt governance is updated as required. The AI Systems Register references model families rather than specific version strings to reduce administrative burden from routine model updates.

Model Limitations: AI outputs may be incomplete or omit relevant site-specific constraints due to limitations in underlying datasets and model interpretation. This risk is mitigated through the limitations disclosure in every GeoScreen report and the requirement for professional review where outputs inform consequential decisions.

8.0 HUMAN OVERSIGHT AND ACCOUNTABILITY

- AI systems used by Magnum GSI Ltd do not make autonomous decisions. All outputs are subject to human interpretation and professional judgement prior to use.
- All AI-generated reports include a disclosure statement confirming the use of AI in their preparation, in accordance with the Limitations section of the GeoScreen report templates.
- GeoScreen reports include a limitations section advising that outputs should be reviewed by a qualified geotechnical engineer prior to engineering or planning reliance.
- Human review must be proportionate to risk and complexity. Where outputs inform engineering, planning, or financial decisions, review by a suitably qualified geotechnical professional is required.
- No warranty is given as to the completeness or accuracy of AI-generated outputs. Outputs are indicative only and should not be relied upon without appropriate professional review.
- Magnum GSI Ltd retains full professional responsibility for the outputs of AI systems used in its services, irrespective of the extent of AI involvement in generating those outputs.
- AI outputs that may have legal, safety, or significant financial consequences are subject to review by a Chartered Geotechnical Engineer before being professionally relied upon.
- Clients are advised in writing that GeoScreen Tier 0 and Tier 1 reports are automated desk-based screening reports and are not a substitute for a professionally authored report or intrusive ground investigation where site conditions require it.
- Clients are responsible for ensuring they understand the limitations of AI-assisted outputs prior to reliance.

9.0 TRANSPARENCY AND DISCLOSURE

Magnum GSI Ltd is transparent about its use of AI in service delivery. All disclosures are presented in plain language suitable for non-technical users:

- All client-facing reports generated using AI include a disclosure statement in the Limitations section of the GeoScreen report, confirming AI-assisted generation and the basis for professional review.
- This AI Systems Management Policy is published on the Magnum GSI Ltd website.
- Clients may request further information about the AI systems used in the preparation of their reports by contacting Enquiries@magnumgsi.com.
- Magnum GSI Ltd will notify clients of any material change to the AI systems used in the preparation of their reports.
- GeoScreen is disclosed to clients as a controlled, rule-governed geotechnical decision-support product, not a general-purpose AI content generator.
- The limitations of AI outputs, including the possibility that outputs may be incomplete or omit site-specific constraints, are disclosed to clients in every report.

10.0 RESPONSIBILITIES

Corporate Governance — The Director of Magnum GSI Ltd (Tristan Morgan) holds overall responsibility for AI governance, including approval of new AI systems, annual review of this policy, approval of the AI Risk Register, and sign-off on the AI Systems Register.

Technical Management — Responsible for maintaining AI system integrations, API credential security, prompt governance rules, quality assurance testing, and maintenance of the AI Risk Register.

Employees — Required to use AI tools only in accordance with this policy, to report concerns or anomalous outputs to the Director, and not to submit client data to personal or unapproved AI tools or services.

Freelancers and Subconsultants — Must not use client or project data in personal AI tools or accounts without prior written approval from the Director. Must comply with this policy and with PO7.00 and PO9.00.

11.0 COMMUNICATION OF THE POLICY

This policy will be displayed on the Magnum GSI Ltd website and communicated to all employees, freelancers, and subconsultants during induction and as required. It will be reviewed at least annually and updated following any material change to Magnum GSI Ltd AI systems or relevant legislation.

12.0 REVISION HISTORY

Version	Date	Description of change
1.0 / 1.2	June 2026	First issue; ISO 42001 and ISO 27701 alignment and QA corrections
2.0	31 October 2026	Annual review. Header title corrected (it read 'GeoScreen Website Privacy Policy'). AI outputs redefined: hazard ratings are set by explicit rules, not AI; AI writes the narrative, reads historical maps (two readings must agree) and runs the pre-issue QA check. Data minimisation wording made exact. QA gate added. Limitations section references corrected for Tier 0 and Tier 1.

Signed by.



Tristan Morgan

Director.

Date: 31 October 2026