Mine Closure 2025 – S Knutsson, AB Fourie & M Tibbett (eds)

doi:10.36487/ACG_repo/2515_66

Generating a Global Industry Standard on Tailings Management knowledge base: harnessing artificial intelligence for enhanced decision-making in tailings management

Simon Gibbons [a,*], Charissa Tomlin [b]

[a] *ERM, UK*

[b] *ERM, Australia*

Abstract

The modern world is accustomed to seeing news articles either praising the use of artificial intelligence (AI) or lamenting the dangers it poses to humanity. Large corporations integrate AI into workflows to facilitate faster and more in-depth analyses of vast data sets. But how does general AI fit into everyday work lives? Is one confined to asking generic or censored questions to public AI tools, or, if fortunate, accessing some form of in-house AI support? The good news is that AI use is not limited, and a multi-million-dollar implementation is not needed to embark on this journey.

This paper explores how AI tools have been utilised, and specifically how it can support a range of requirements under the Global Industry Standard on Tailings Management (GISTM) by helping collate and curate knowledge on tailings, as well as reviewing archived documents and studies. It also investigates how that information can be interrogated, free from the realm of misinformation on the wider web, as well as provide an analysis of how portfolios of sites can share data and, when necessary, make comparisons with publicly available information. The intent of these tools is to enable gathering of information that aids in making better decisions and identifying areas requiring deeper understanding.

Keywords: *GISTM, tailings, knowledge base, artificial intelligence, mining*

1 Introduction

The rapid evolution of artificial intelligence (AI) has dominated both technical discourse and mainstream media in recent years. Headlines oscillate between portraying AI as humanity's greatest achievement or its most significant threat, creating a narrative disconnect between public perception and practical implementation. While large technology corporations and financial institutions deploy sophisticated AI systems to analyse vast datasets and optimise operations, many industries – particularly those with specialised technical requirements like mining – struggle to bridge the gap between AI's theoretical potential and tangible operational benefits.

Mining operations worldwide face increasing regulatory scrutiny, with tailings management emerging as a critical focus area following catastrophic failures like those at Mount Polley (Canada, 2014), Fundão (Brazil 2015) and Brumadinho (Brazil 2019). The Global Industry Standard on Tailings Management (GISTM), launched in August 2020 (International Council on Mining and Metals et al. 2020), represents the industry's response to these incidents (noting that it was originally the investor community, led by the Church of England Pension Board who demanded improvement in industry performance and disclosure), establishing rigorous requirements for the entire tailings facility life cycle. However, implementing these standards

* Corresponding author. Email address: simon.gibbons@erm.com

presents significant challenges, particularly in knowledge management, historical data integration, and consistent cross-site analysis – challenges that well-deployed AI solutions are uniquely positioned to address.

Despite the mining industry's traditional reputation for technological conservatism, the sector has increasingly embraced digitalisation initiatives, with varying degrees of success. Large mining corporations have invested in predictive maintenance, autonomous equipment, and advanced sensing technologies, but these implementations often represent isolated applications rather than comprehensive digital transformation. More importantly, these initiatives typically demand substantial capital investment, specialised expertise, and lengthy implementation timelines – resources that may be unavailable to mid-tier producers or operations in developing regions.

The knowledge management requirements specified by GISTM present a compelling opportunity for targeted AI application. The standard explicitly requires operators to maintain comprehensive records, conduct regular assessments, and demonstrate thorough understanding of a facility's construction and operational history. For legacy sites with decades of operation, this often means accessing, organising, and interpreting thousands of documents across multiple formats – from scanned paper records to digital reports and monitoring data. Traditional approaches to this challenge involve teams of engineers and technicians manually reviewing documents, a process that is not only resource-intensive but also prone to oversight and inconsistency.

Contemporary AI tools, particularly those leveraging natural language processing (NLP) and machine learning, offer an alternative approach. These technologies can process vast document repositories, extract relevant information, identify patterns, and organise knowledge in accessible formats. Importantly, recent developments in foundation models and specialised AI applications have dramatically reduced the technical and financial barriers to implementation. Operations can now leverage these capabilities without multi-million-dollar investments or dedicated data science teams, democratising access to powerful analytical tools as illustrated in Figure 1 below.

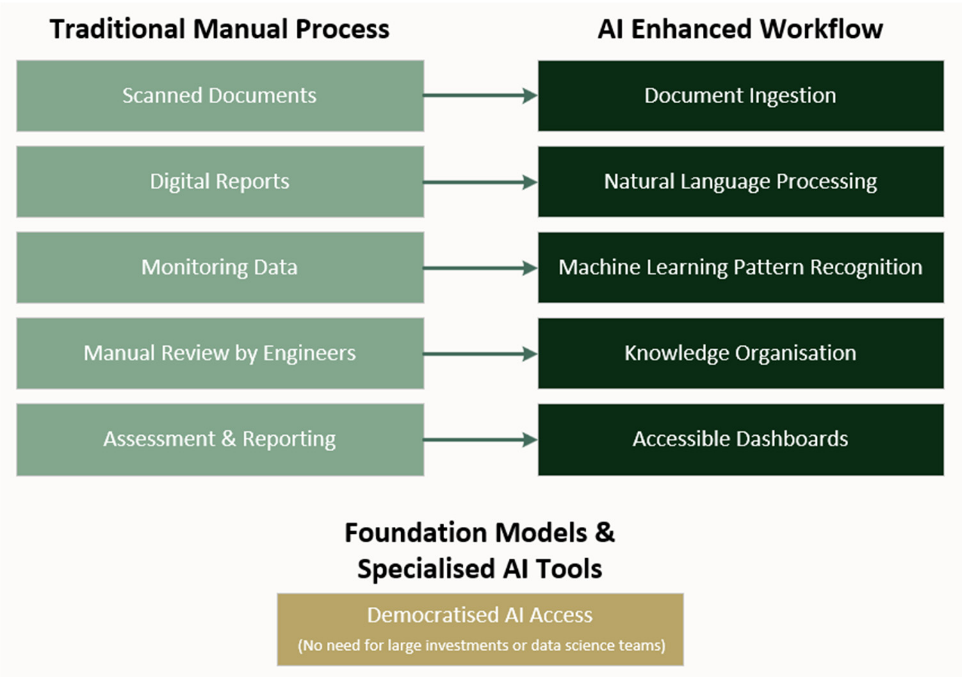

Figure 1 Artificial intelligence enhanced knowledge management for legacy sites under GISTM

This democratisation coincides with growing recognition that effective tailings management requires both technical excellence and comprehensive knowledge systems. The GISTM emphasises the importance of maintaining “knowledge about the tailings facility that is accessible and comprehensive” (Topic 4, Requirement 4.1) and stipulates that “new data and information [should be] taken into account” throughout a facility’s life cycle (Requirement 7.3), (International Council on Mining and Metals et al. 2020). Meeting these requirements demands not just sophisticated data handling, but intelligent information synthesis—precisely the capability that modern AI tools can provide. Figure 2 illustrates the six key topics within GISTM and the specific areas where knowledge management plays a core role in understanding and managing risk.

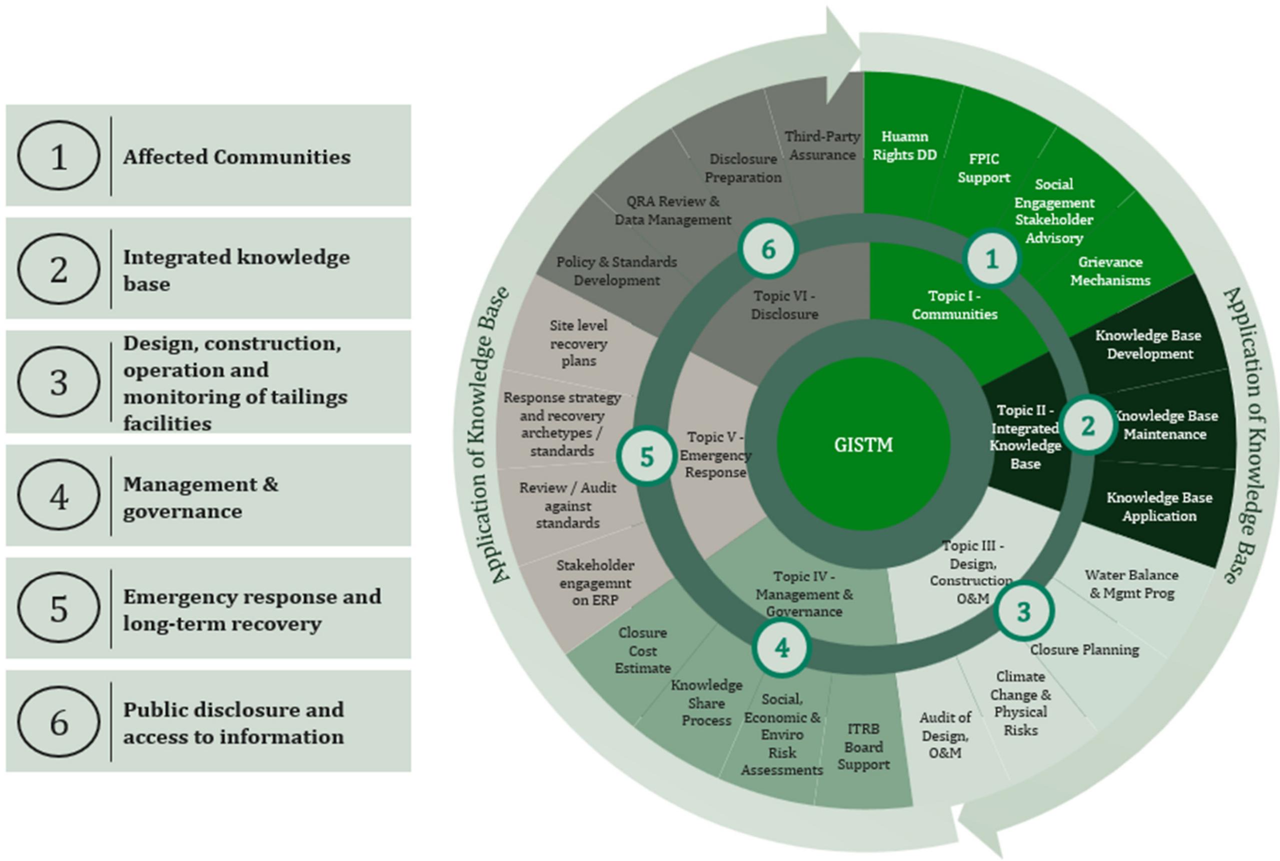

Figure 2 Global Industry Standard on Tailings Management topics

This paper explores practical applications of accessible AI technologies in supporting GISTM implementation, with specific focus on knowledge collation, document analysis, information interrogation, and portfolio management. Rather than theoretical possibilities, we examine implemented solutions that demonstrate how mining operations of various sizes can leverage AI capabilities to enhance tailings management without prohibitive investment. Through case studies and practical examples, we illustrate how these tools enable better decision-making and identify areas requiring deeper understanding – key objectives in the ongoing effort to improve tailings management practices globally.

2 Background and literature review

GISTM emerged as a direct response to catastrophic tailings dam failures that resulted in significant loss of life, environmental damage, and harm to the mining industry's reputation. Developed through a multi-stakeholder process involving the International Council on Mining and Metals (ICMM), the United Nations Environment Programme (UNEP), and the Principles for Responsible Investment (PRI), the standard establishes 15 principles and 77 auditable requirements across six topic areas. Knowledge management underpins many of these requirements, with explicit mention throughout the standard.

Requirement 4.1 stipulates that operators "develop and maintain knowledge about the tailings facility and its surroundings that is accessible and comprehensive," while Requirement 7.8 demands "a system to manage the quality and adequacy of the tailings facility data.", (International Council on Mining and Metals et al. 2020). These and other requirements create substantial information management challenges, particularly for facilities with decades of operational history or those with incomplete records following ownership transitions.

A study by Schoenberger (2016) examined knowledge management frameworks for tailings facilities that align with what would later become GISTM requirements, focusing on methods for maintaining comprehensive records and facilitating knowledge transfer between operational phases. Similarly, following the Fundão Tailings dam disaster, Morgenstern et al. (2016) prepared a highly influential report, which helped shape GISTM development, specifically addressing the knowledge management failures which contributed to the disaster and the associated response, and which recommended a structured approach to information management for tailings facilities, and which directly informed the GISTM requirements.

Barros et al. (2020) further discusses integrated management systems that incorporate knowledge management practices aligned with the emerging GISTM requirements, particularly focusing on methods for information integration and accessibility.

Existing approaches to tailings knowledge management range from basic document repositories to sophisticated integrated data management systems. Traditional methods typically involve manual document review, spreadsheet-based tracking, and fragmented databases that struggle to provide comprehensive, accessible information. Goebel and Fülöp (2020) documented that mining engineers spend an average of 7.5 hours per week searching for information across disparate systems, highlighting the inefficiency of conventional approaches. A 2021 industry guide on GISTM (Chiesa et al. 2021) specifically addresses the implementation challenges of GISTM and presents a substantial discussion of knowledge management requirements and the practical difficulties companies face in meeting them.

The application of AI in mining has gained momentum in recent years, though primarily focused on operational optimisation rather than knowledge management. A comprehensive review by Barnewold & Lottermoser (2020) and Durant-Whyte et al. (2015) each identified predictive maintenance, autonomous equipment, and process optimisation as the dominant AI applications in mining, with knowledge management representing less than 8% of documented implementations. This imbalance presents both a gap in the literature and an opportunity for focused development.

Research on AI for document analysis and knowledge extraction in other regulated industries offers transferable insights. The legal sector has embraced natural language processing for contract analysis and case research (Chalkidis & Kampas 2019), while the pharmaceutical industry utilises similar techniques for literature review and regulatory compliance (Johnson et al. 2021). These applications demonstrate the potential for domain-specific AI tools that combine general capabilities with industry-specific knowledge. An illustration of the workflow for such an approach is shown in Figure 3.

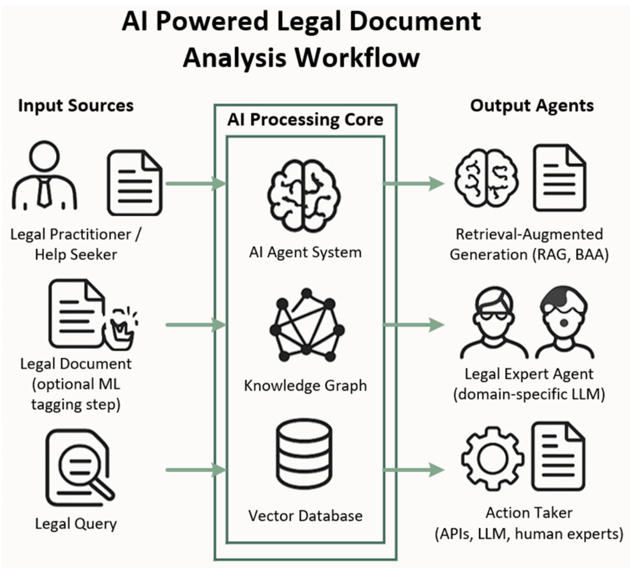

Figure 3 Workflow diagram of artificial intelligence powered legal document review

Recent advances in foundation models like GPT-4 (OpenAI 2023), Claude (Anthropic 2023), and LLaMA 2 (Touvron et al. 2023) have dramatically reduced implementation barriers for text-based AI applications. These models demonstrate remarkable capabilities in understanding context, extracting information, and generating coherent summaries across diverse document types. As noted by Bommasani et al. (2022), these models "enable adaptation to specialised domains with minimal technical expertise," a critical consideration for mining operations with limited data science resources.

Bridging the gap between AI's potential and real-world impact remains a challenge. Chui et al. (2018) highlight that advanced deep learning techniques have the potential to generate trillions of dollars in value across industries, but capturing this value requires overcoming significant technical and organisational hurdles. This underscores the importance of translating AI's theoretical capabilities into practical implementations in mining.

The intersection of these developments – GISTM's knowledge requirements, mining's digitalisation journey, and AI's increasing accessibility – creates a compelling case for systematic exploration of AI-enabled knowledge management for tailings facilities. The remainder of this paper examines this intersection through implemented case studies that span the continuum from simple support tools to comprehensive organisational systems.

3 Methodology: artificial intelligence tools selection and implementation

Identifying and implementing appropriate AI tools for tailings management requires a structured approach that balances technical capabilities with practical constraints. This section outlines the methodology used to select, customise, and deploy AI solutions across multiple mining operations, ranging from single-site implementations to enterprise-scale systems.

3.1 Evaluation framework

The selection of AI tools follows a four-dimensional evaluation framework designed to prioritise practical implementation over theoretical capabilities:

1. Accessibility: assessed by technical implementation requirements, necessary expertise, and integration complexity with existing systems. Tools requiring minimal specialised knowledge score higher in this dimension.
2. Adaptability: measured by the ability to customise for mining-specific terminology, incorporate domain knowledge, and evolve with changing requirements. Foundation models with fine-tuning capabilities typically excel here.
3. Scalability: evaluated through performance with increasing document volume, user base expansion capabilities, and cross-site implementation potential. Both technical scalability and licensing/cost scalability are considered.
4. Security and governance: rated based on data sovereignty options, permission controls, and alignment with corporate information security policies. This dimension is particularly important for sensitive operational data.

Each potential solution can be scored across these dimensions using a weighted scoring system that prioritised dimensions according to organisational context – smaller operations typically weighted accessibility higher, while multinational corporations emphasised security and governance.

3.2 Technology components

The implemented solutions incorporated various technologies tailored to specific requirements:

1. Document processing pipeline: utilised optical character recognition (OCR) technologies (primarily Tesseract and ABBYY FineReader) for converting scanned documents, combined with custom pre-processing to handle mining-specific documentation formats.
2. Foundation models: employed large language models (LLMs) including GPT-4, Claude, and open-source alternatives (LLaMA 2, Falcon) depending on security requirements and budget constraints. Private deployment options were evaluated for sensitive information.
3. Vector databases: implemented Pinecone, Weaviate, or Qdrant to enable semantic search capabilities across document repositories, with embedding models selected for compatibility with mining technical documentation.
4. Interface technologies: developed through a combination of existing enterprise platforms (SharePoint, Microsoft Teams) and custom web interfaces, prioritising integration with established workflows.
5. Data connectors: created for common mining software packages including Deswik, Maptek, and Leapfrog to enable bi-directional information flow between knowledge systems and operational tools.

3.3 Data integration

In addition to the data processing steps mentioned above, there is an opportunity to integrate high-resolution data collection tools with AI, which can help make sense of the data and identify patterns within it. Although this approach is not included in the case studies, it is important to document it here when considering future applications.

- AI with live data streams: beyond historical documents, tailings facilities continuously generate data – e.g. sensor readings (piezometer pressures, inclinometer movements), drone and satellite imagery, remote sensing data, etc. AI could correlate these live data feeds with the historical knowledge. For example, an AI system might detect an anomaly in sensor data and cross-reference it with past incident reports or design changes in the knowledge base to warn of potential issues earlier. This proactive, combined use of current and historical data would greatly enhance risk management (essentially moving from reactive analysis to predictive monitoring).
- Spatial and imagery analysis: much of tailings risk assessment involves visual and spatial data. AI-driven image analysis (computer vision models) could scan aerial photographs or satellite images for changes in a dam’s appearance (e.g. new seepage spots or slope bulges) over time. Those findings, when linked to a text-based knowledge base, provide a more holistic view.
- Samarco failure analysis: the forensic investigation of the Fundão dam failure relied on aerial imagery and piezometer data over time. An AI system could potentially have traced such subtle warning signs by combining spatial data trends with the historical design and inspection records, potentially flagging the issue in advance. This underscores the value of merging real-time monitoring with historical knowledge.

3.4 Implementation approach

Implementation followed a phased methodology designed to minimise disruption while demonstrating incremental value:

- Phase 1: Document Digitisation and Organisation**.** This initial phase focused on creating a machine-readable set of tailings documentation through scanning, OCR processing, and basic metadata extraction. Documents were categorised using a standardised taxonomy aligned with

GISTM requirements, enabling basic search and retrieval. During this phase data cleansing approaches are deployed to remove out of date or unreliable information – marking data gaps as appropriate. Significant effort can be required for data cleaning and standardisation: e.g. resolving inconsistent units, typos, or formatting in decades-old reports, and standardising terminology (tailings documents can vary greatly in language). Automating this (where possible) frees up specialists' time for higher-level analysis.

- Phase 2: Knowledge Extraction and Enhancement. The second phase employed AI tools to extract key information from processed documents, including design parameters, inspection results, and monitoring data. This information was structured according to facility-specific information models and enhanced with cross-references to source documents.
- Phase 3: Interface Development and Integration. User interfaces were developed to provide role-specific access to the knowledge base, with separate views for operators, engineers, managers, and auditors.
- Phase 4: Advanced Analytics and Decision Support. The final phase focused on developing analytical capabilities, including trend analysis, anomaly detection, and comparison against design criteria. These capabilities transformed the system from a knowledge repository to a decision support tool.

3.5 Evaluation metrics

The effectiveness of implemented solutions was measured using both quantitative and qualitative metrics:

1. Information retrieval efficiency: measured by time required to locate specific information compared to pre-implementation baseline.
2. Knowledge extraction accuracy: evaluated through manual verification of AI-extracted parameters against source documents.
3. User adoption: tracked through system usage statistics and structured feedback sessions.
4. GISTM compliance support: assessed by mapping system capabilities to specific GISTM requirements and evaluating coverage as well as through audit feedback.

These metrics were tracked throughout implementation to guide system refinement and document achieved benefits for organisational stakeholders.

4 Case studies: artificial intelligence applications in Global Industry Standard on Tailings Management compliance

The following case studies are a combination of potential use case examples and real-world deployments of technology to build knowledge base systems from the authors own experience. Quantification of benefits is based on approximations from various AI implementation studies across industry. For example, McKinsey & Company (2023), document how organisations implementing AI solutions have seen 20-40% increases in workforce productivity and 30-50% reductions in time spent on information discovery and retrieval tasks. Davenport & Ronanki (2018) report on survey results showing that organisations implementing document intelligence and natural language processing solutions achieved 50–75% reductions in time spent on document review and information extraction.

Further studies by Vergara et al. (2021), found that digital twins and other digital technology implementation, including document management systems, achieved a 30–40% reduction in information retrieval time and greater degrees of documentation compliance. A further study by Jha et al. (2022) provides specific metrics on efficiency improvements from AI-based document processing in mining contexts, reporting 22–38% reductions in document processing times and 28–42% improvements in knowledge retrieval accuracy. Similar performance outcomes were also reported by Vidal-Silv et al. (2021) with a specific focus on knowledge management in the mining sector.

4.1 Case study 1: knowledge collation and curation

4.1.1 Background

A mid-tier gold producer with operations across three continents faced significant challenges preparing for GISTM conformance audits. The company's oldest site had accumulated over 40 years of tailings documentation across various media – from hand-drawn designs on microfiche to modern digital instrumentation data. Critical knowledge existed in silos, with design information maintained by engineering, operational data held by production teams, and monitoring results managed by environmental departments.

4.1.2 Implementation

A two-tier AI approach was implemented to address these challenges.

- First, a document processing pipeline was established to digitise all historical documentation. This system employed OCR to generate searchable text documents. These documents underwent semi-automated classification using a custom-trained model based on the open-source LLaMA model that categorised content according to GISTM-aligned taxonomy.
- Second, a knowledge extraction system was deployed using a fine-tuned large language model (based on LLaMA 2). This system analysed processed documents to identify key parameters, including:
  - design specifications and construction details
  - material properties and deposition methods
  - monitoring instrumentation locations and specifications
  - historical inspection findings and recommendations
  - remediation activities and modifications.

The extracted information populated a structured knowledge graph that represented the complex relationships between physical components, monitoring activities, and governance processes. This graph served as the foundation for a comprehensive "facility knowledge base" as required by GISTM.

A knowledge graph (see Figure 4) is a structured representation of information that captures relationships between entities – such as people, places, concepts, or documents – in a graph format. It's designed to enable machines (and humans) to understand and reason about data more effectively.

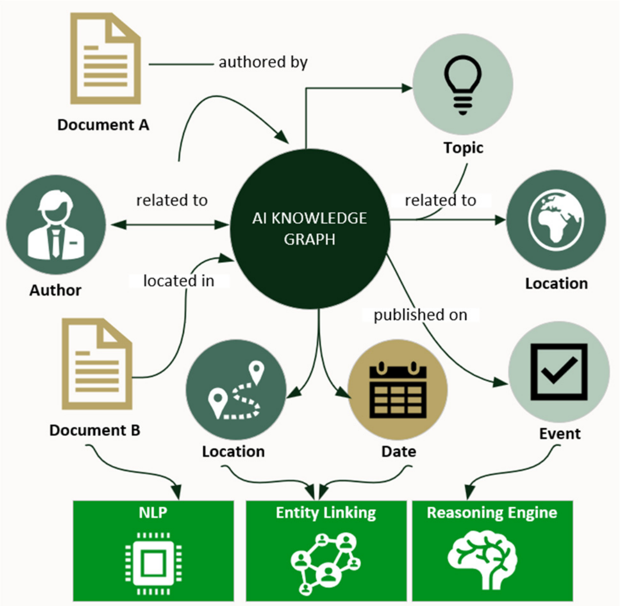

Figure 4 Knowledge graph illustration

Key components:

- nodes: represent entities (e.g. “Document A”, “AI Model”, “Author”)
- edges: represent relationships between entities (e.g. “authored by”, “references”, “is part of”)
- attributes: additional metadata about nodes or edges (e.g. publication date, confidence score).

In the context of AI and document analysis a knowledge graph can be used to:

- link extracted facts from documents (e.g. “Company X acquired Company Y in 2023”)
- connect related concepts across multiple sources
- enable semantic search and reasoning (e.g. “Find all documents related to AI regulations authored by UK-based researchers”).

4.1.3 Results

The implemented system significantly reduced information retrieval time compared to manual methods, with engineers reporting that queries that previously required days of document review could be completed in under an hour.

Perhaps most significantly, the system identified previously unknown information gaps, particularly regarding construction details from facility expansions in the 1990s. These gaps were subsequently addressed through geophysical investigations, significantly improving the operation's understanding of actual facility conditions.

4.2 Case study 2: document review and analysis

4.2.1 Background

A multinational mining corporation undertaking GISTM implementation across 27 active tailings facilities faced the challenge of reviewing thousands of historical inspection reports, consultant recommendations, and audit findings. It was estimated that comprehensive manual review would require more than 8,000 person-hours (approximately 1 hour per document) – resources that would divert environmental and social team members from current operational requirements.

4.2.2 Implementation

An AI-powered document analysis system was implemented to streamline this review process. The system combined NLP with mining-specific pattern recognition to:

1. identify recommendations, findings, and action items across diverse document formats
2. track the status of identified actions across subsequent documentation
3. flag potentially unresolved or recurring issues
4. detect patterns in inspection findings that might indicate progressive deterioration
5. generate comprehensive action registers with supporting evidence.

The system was deployed using a secure cloud infrastructure with role-based access controls that maintained document confidentiality while enabling cross-site pattern recognition. Corporate governance teams maintained oversight through a dashboard that visualised compliance status and outstanding actions across the portfolio.

4.2.3 Results

The AI-powered review process not only accelerated document analysis but also enabled multi-dimensional pattern recognition across historical inspection reports, monitoring data, and incident logs. The system employed NLP specifically to extract and cluster key terms such as “seepage,” “toe,” “piezometric rise,” and “drainage inefficiency” from unstructured text. These keyword clusters were cross-referenced with structured datasets including piezometric pressure readings, water quality indicators (e.g. turbidity, pH, conductivity), and visual inspection logs tagged with geospatial metadata.

Machine learning models were trained to detect temporal and spatial correlations between these indicators. For example, recurring mentions of “toe wetness” in inspection narratives were found to align with anomalous pore pressure trends and elevated turbidity levels downstream – patterns that had not previously been flagged as systemic due to their distribution across facilities and timeframes. Visual data, including drone imagery and satellite-based Interferometric Synthetic Aperture Radar (InSAR) deformation maps, were also integrated to validate surface expression of seepage zones. A simple illustration of the workflow is shown in Figure 5.

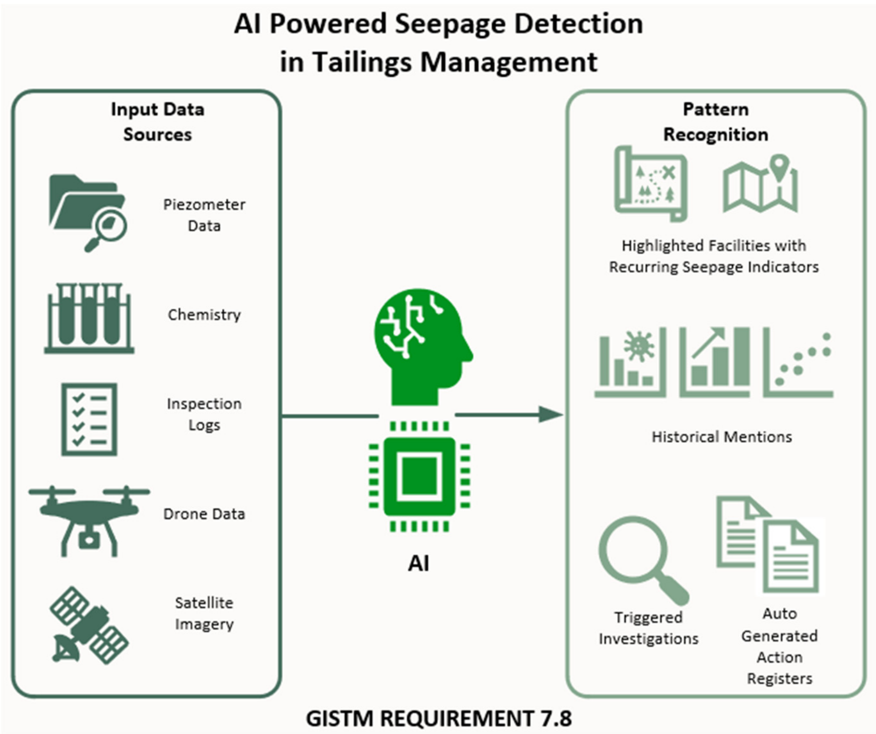

Figure 5 Artificial intelligence powered pattern recognition

This multi-source synthesis enabled the AI system to highlight three facilities where seepage-related indicators had been repeatedly noted but not escalated. The discovery prompted a targeted engineering review, which confirmed the need for drainage upgrades and led to the implementation of corrective works. These findings were captured in automatically generated action registers, directly supporting GISTM Requirement 7.8 by providing traceable evidence of issue identification, analysis, and resolution.

4.3 Case study 3: information interrogation systems

4.3.1 Background

A sustainability consultancy supporting multiple clients with GISTM implementation sought to leverage its collective expertise while maintaining strict client confidentiality. The firm's specialists possessed valuable insights from diverse implementations but struggled to effectively share this knowledge across projects due to information barriers and time constraints.

4.3.2 Implementation

The consultancy developed an AI-powered information interrogation system using a combination of open-source tools including LLaMA and Microsoft Copilot Studio that combined:

1. a private knowledge base containing anonymised lessons learned, technical approaches, and implementation strategies derived from client engagements
2. a customised large language model fine-tuned on tailings management literature, standards, and guidelines
3. a rigorous fact-checking mechanism that verified generated responses against source documentation.

An illustration of the project setup page and an interactive chat with the AI agent on the specifics of monitoring data (based solely on the agents education) are provided as illustrations in Figure 6.

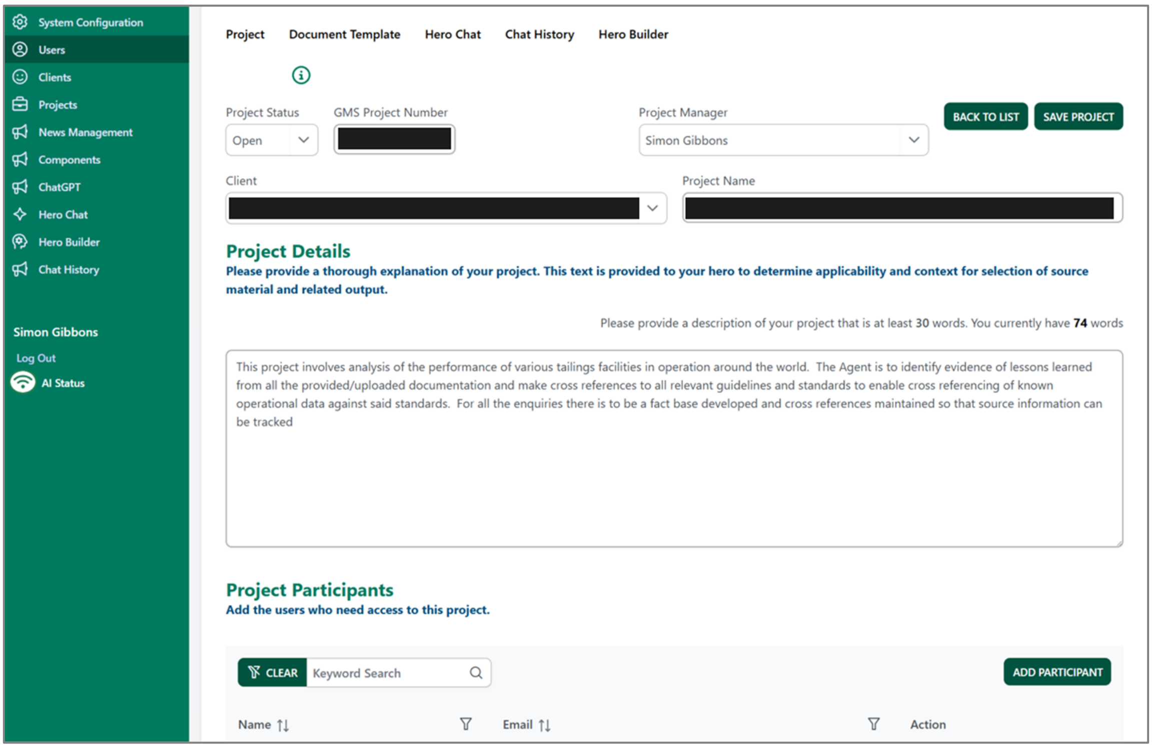

(a)

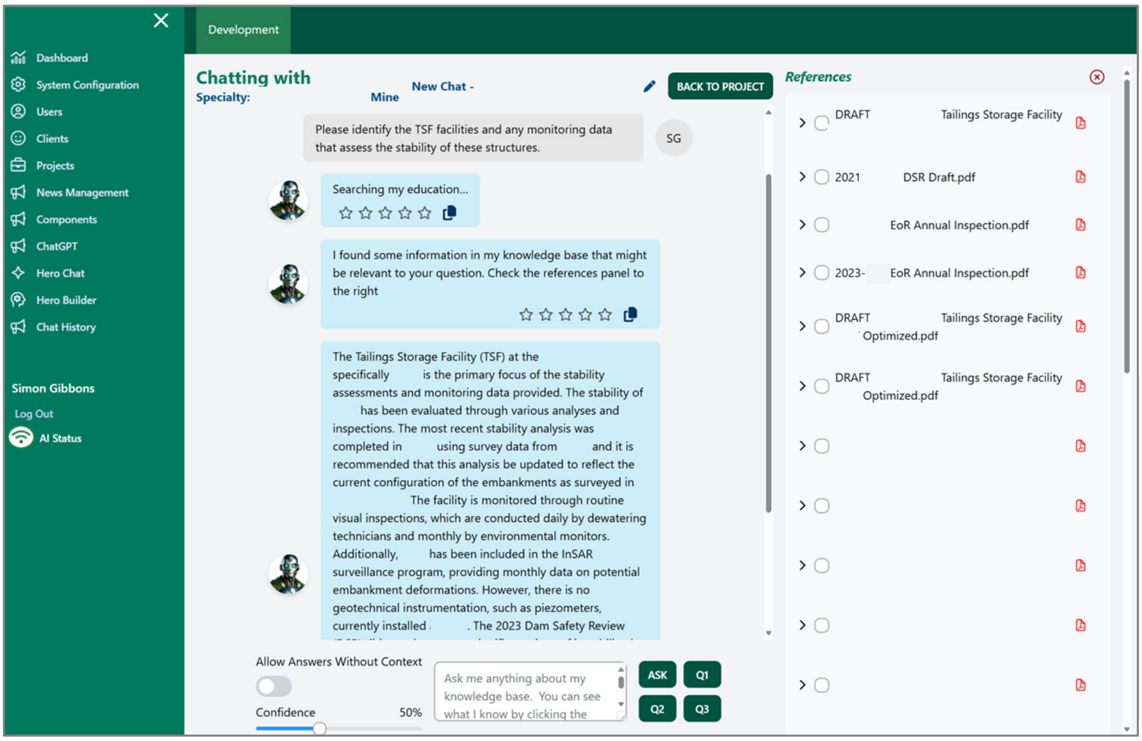

(b)

Figure 6 Artificial intelligence development for private knowledge base setup (a) and AI agent engagement (b)

The system enabled consultants to query complex tailings management questions and receive evidence-based responses that cited relevant precedents, applicable standards, and technical considerations. Unlike public AI tools, this system operated within a controlled information environment free from potential misinformation and focused exclusively on verified tailings knowledge.

4.4 Case study 4: portfolio management and benchmarking

4.4.1 Background

A diversified resources company operating 35 tailings facilities across four commodities (copper, zinc, nickel, and coal) faced challenges in standardising its approach to GISTM implementation. Each commodity business maintained separate technical teams with distinct approaches to tailings management, complicating corporate governance and comparative performance assessment.

4.4.2 Implementation

An enterprise-scale AI system was implemented to enable portfolio-wide knowledge management while accommodating site-specific requirements. The system featured:

1. a centralised document repository with federated access controls that maintained business unit autonomy while enabling corporate oversight
2. automated extraction of key performance indicators aligned with GISTM requirements
3. standardised terminology mapping that preserved site-specific nomenclature while enabling cross-site comparison
4. benchmarking capabilities that identified best practices across the portfolio
5. integration with public databases of environmental and social data to provide external context.

The portfolio management system has enabled the first comprehensive cross-commodity analysis of tailings performance within the organisation, revealing significant variations in monitoring practices, risk assessment methodologies, and governance approaches. These insights prompted a corporate initiative to harmonise critical practices while maintaining necessary site-specific adaptations. Figure 7 shows an example of the AI powered document review and data extraction system.

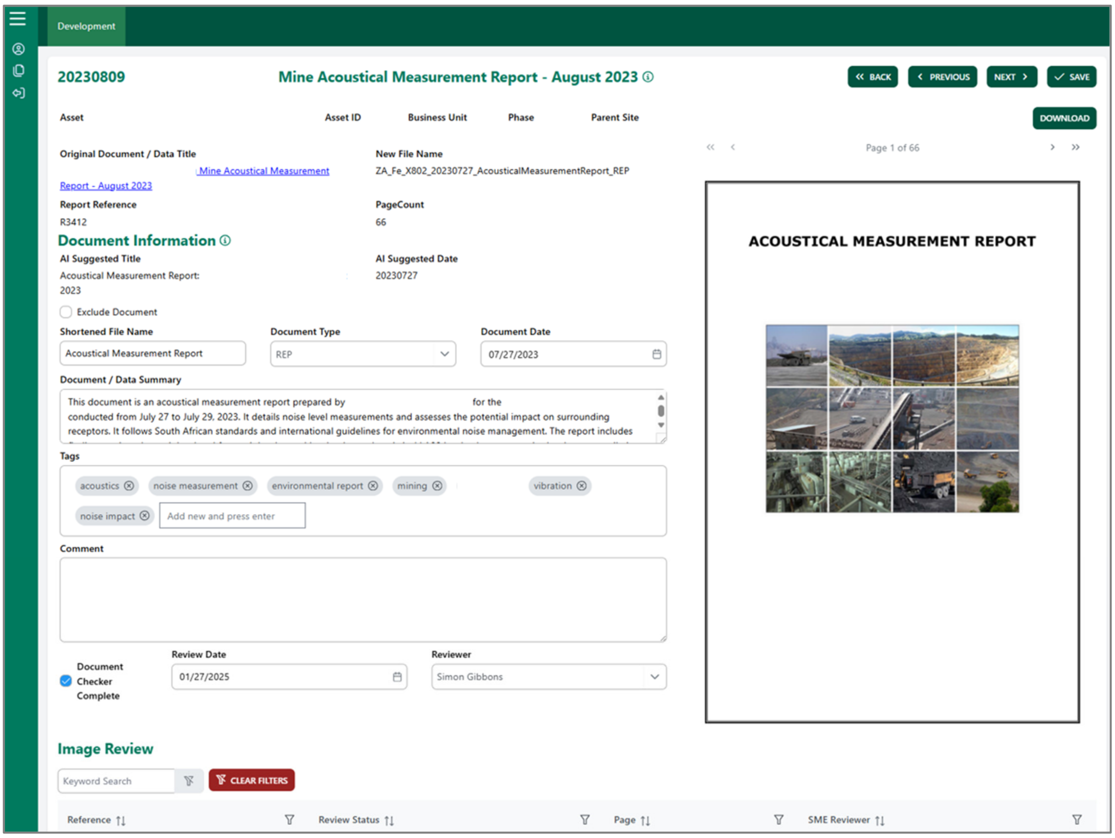

Figure 7 Artificial intelligence document review and knowledge extraction application

Comparative analysis identified that nickel operations consistently achieved more rapid implementation of instrumentation recommendations, leading to an investigation that revealed effective procurement practices subsequently adopted across the organisation. Conversely, the analysis identified that coal operations had developed superior community engagement practices that were subsequently shared with other commodities.

5 Results and discussion

The implementation of AI-enabled knowledge management systems across diverse mining operations has, with varying implications depending on organisational context and implementation approach, demonstrated consistent benefits for GISTM implementation. This section synthesises key findings and discusses their broader implications for the mining industry:

5.1 Performance improvements

Quantitative analysis across implementations reveals several consistent performance improvements:

- Time efficiency: AI-enabled systems reduced information retrieval time by up to 90% compared to manual methods, with the greatest improvements observed for complex queries spanning multiple document types
- Knowledge quality: the systematic extraction and organisation of information significantly improved knowledge completeness compared to pre-implementation baselines. More significantly,

AI systems identified previously unknown information gaps in most of the implementations, enabling targeted investigations to address these deficiencies

- Governance effectiveness: organisations implementing portfolio-level systems reported improvement in consistency of GISTM interpretation across operations and a significant reduction in time required to prepare consolidated governance reports
- Cross-learning effectiveness: operations utilising AI-enabled knowledge sharing documented more instances of cross-site learning than those using conventional approaches, with a higher implementation rate for identified best practices. These improvements have contributed to more robust tailings management practices and enhanced readiness for GISTM conformance audits, with early adopters reporting significantly fewer major findings during independent reviews
- Cost-benefit analysis: various studies have outlined ways to calculate the costs and benefits of knowledge management systems. For example, Samimi and Najafzadeh (2021) present a quantitative approach to evaluate return on investment (ROI) for digital technologies in mining (including knowledge management), providing empirical cost data across different operation sizes. Similarly, Mikalef et al. (2018) note that implementation costs for knowledge platforms can range widely – from about USD 75k–150k for departmental-level systems up to USD 500k–1.5 million for enterprise-scale platforms. These figures, coupled with the rapid advancement of AI since 2018, suggest that modern solutions may achieve lower costs and faster payback periods. The analyses consistently show favourable economics: reduced labour hours (immediate cost savings)
- Direct cost savings: a study by Ghasemaghaei & Calic (2020) found that reduced professional time for information retrieval represented immediate benefits, with documented annual savings of USD 50,000-280,000 depending on organisation size and information complexity. Anecdotal information from the presented case studies implies individual site savings of approximately USD 25,000 per site solely based on data retrieval times from implemented systems
- Risk reduction value: research undertaken by Katic and Majstorović (2022) quantified risk reduction benefits in high-consequence industries (including mining), finding risk exposure reductions valued between USD 300,000–2.1 million annually per operational site when using comprehensive knowledge management systems. Whilst the studies by the author have not specifically demonstrated risk reduction, the fact that the use of AI is highlighting missed recommendations for improvements and systematic challenges at facilities, it could be assumed that the implementation has helped avoid at least one significant near-miss which suggests value of risk reduction could readily reach the million dollar plus range
- Implementation timeframes: according to Samimi & Najafzadeh (2021) payback periods for knowledge management systems in mining operations typically range from 10–14 months for targeted implementations to 18–36 months for enterprise-wide platforms. Evidence from the authors own experience suggests this is plausible and with potential for the recent burst in AI development potentially providing more rapid implementation timeframes – potentially down to a few weeks in the instance where organisations are choosing to deploy AI assistants throughout their business.

5.2 Ongoing challenges

Despite demonstrated benefits, several consistent challenges emerged across implementations:

- Data quality limitations: historical documentation often contained inconsistencies, contradictions, and errors that complicated automated processing. Successful implementations incorporated uncertainty quantification and human verification mechanisms rather than attempting fully automated solutions

- Expertise integration: effective systems required integration of domain expertise throughout development, particularly for knowledge extraction rules and evaluation criteria. Organisations that treated AI implementation as primarily a technical IT challenge achieved significantly poorer results than those with multidisciplinary implementation teams
- Change management: user adoption presented consistent challenges, particularly among experienced personnel accustomed to established methods. Implementation success correlated strongly with comprehensive change management programs that demonstrated immediate value to end users. Whilst AI is intuitive it also takes specific skills and training to use natural language processing to its maximum potential. Coaching is required as is cross-verification of results
- Appropriate trust calibration: both over-reliance and under-utilisation of AI systems reduced effectiveness. Successful implementations established clear guidelines for appropriate system use, including explicit identification of limitations and verification requirements
- Decision responsibility: all successful implementations maintained clear human responsibility for decision-making, using AI systems as decision support rather than autonomous agents. Governance frameworks explicitly documented this distinction, particularly for high-consequence decisions
- Transparency and explainability: systems varied in their explainability, with simple document retrieval offering clear provenance while complex pattern recognition proved more challenging to explain. Organisations increasingly prioritised explainable AI approaches, particularly for applications directly supporting consequence classification or risk assessment
- Knowledge preservation versus innovation: over-reliance on historical knowledge risked perpetuating outdated approaches. Effective implementations incorporated mechanisms to identify evolving best practices and challenge historical assumptions when appropriate
- Implementation success: according to the IBM Institute for Business Value (Prabhakar et al. 2021), only about half of AI projects ever progress from pilot to full production, largely because organisations must ground AI initiatives in practical business needs rather than hype (Prabhakar et al. 2021). This insight suggests that mining companies should rethink their approach to AI – focusing on targeted, value-driving applications (like tailings knowledge management) to ensure their AI projects deliver tangible results.

5.3 Practical challenges

There are also additional practical challenges around adoption of AI which include:

- Workforce and culture: resistance from traditional mining professionals who may be wary of AI or comfortable with legacy practices. Buy-in may require training and demonstrating quick wins
- Data bias and quality: AI models can inherit biases from training data or overlook context, which is risky in safety-critical decisions. Therefore, there is a need to carefully curate training data and involve domain experts to check AI outputs
- Trust and transparency: mining teams might trust their own experience over an "AI black box." Providing transparent explanations (e.g. showing which documents or data the AI used to reach a conclusion) helps build user trust
- Change management: introducing AI isn't just a technology upgrade – it requires change management. Companies should encourage a culture where AI is seen as a tool to augment, not replace, human expertise
- Implementation timeline: deployment timeframes across different operation scales can also be a challenge. For instance, a single-site pilot project might get up and running in a few months, whereas a full enterprise-wide rollout (spanning dozens of sites) could take 1–2 years to fully

realise due to integration and training needs. Smaller, focused implementations (targeting specific GISTM requirements) can show results quickly, which can then justify scaling up.

5.4 Governance and ethics

A Deloitte (2022) report emphasises that managing risk and building trust are critical for scaling AI across industries, reinforcing the need for clear ethical guidelines and oversight in AI-driven tailings management. In other words, establishing “trustworthy AI” frameworks and risk controls is just as important as the technology itself when deploying AI in safety-critical domains like mining.

Key areas to address when assessing how to deploy AI include:

- Data security and privacy: there is a need to assure that sensitive operations data is protected – e.g. AI tools are deployed in secure environments, and data sharing follows client agreements
- Transparency and explainability: for safety-critical decisions, the AI’s recommendations need to be explainable. It is important to implement an approach where the AI always provides the source or rationale for any conclusion (like citing the document or data that led to an alert). This makes oversight easier and builds trust
- Human oversight: AI outputs do not override human judgment. For example, an AI-flagged risk must still be evaluated by a qualified engineer (e.g. Engineer of Record for the TSF). Final decisions remain with human experts, with AI acting as a support tool
- Bias and model limitations: there is a risk of algorithmic bias or blind spots – e.g. if the AI was trained mostly on data from large, well-documented sites, it might underperform on smaller or less-documented facilities. This may be mitigated by continuous model evaluation and not using AI in isolation for critical judgments
- Governance policies: ethical AI use needs to follow internal governance standards around its use and these need to be in place ahead of deployment and need to be kept up to date as the technology evolves.

These considerations highlight the importance of thoughtful governance structures around AI implementation, particularly as these technologies become more deeply integrated into tailings management practices.

6 Future directions

The rapid evolution of both AI capabilities and tailings management practices suggests several promising directions for future development in this domain.

6.1 Integration with monitoring systems

Current implementations primarily focus on documentary knowledge, with limited integration of real-time monitoring data. An AI supported knowledge base could interface with geospatial and remote sensing platforms (GIS databases, satellite monitoring systems) and Internet of Things (IoT) sensor networks already in use for dam surveillance. Even if not implemented yet, an AI supported system can enhance rather than replace existing monitoring infrastructure.

Emerging opportunities include:

- automated correlation between observed behaviour and design predictions
- early warning systems that combine historical knowledge with real-time measurements
- dynamic updating of conceptual models based on observed performance
- automated anomaly detection informed by facility-specific design constraints.

These integrations would transform knowledge systems from retrospective repositories to dynamic decision support tools that continuously incorporate new information.

6.2 Advanced multimodal capabilities

Next-generation systems will likely incorporate stronger capabilities for processing diverse information types, including:

- automated interpretation of imagery (satellite, drone, and inspection photographs)
- processing of geophysical data to infer subsurface conditions
- integration of laboratory testing results with field observations
- analysis of operational records to identify depositional patterns.

These capabilities would enable more comprehensive facility understanding, by bridging traditionally separate information streams.

6.3 Collaborative intelligence

Future systems will likely emphasise collaborative intelligence that combines human expertise with AI capabilities:

- interactive exploration of complex facility data
- AI-suggested investigations based on identified knowledge gaps
- simulation tools that allow testing of intervention strategies
- knowledge capture systems that preserve rationale for key decisions.

These approaches would maintain human judgment as the foundation of tailings management while leveraging AI to extend analytical capabilities and institutional memory.

6.4 Regulatory evolution

The relationship between AI-enabled knowledge management and regulatory frameworks will continue to evolve through:

- potential acceptance of AI-assisted analyses in conformance demonstrations
- development of standards for AI application in safety-critical mining contexts
- requirement for transparent, auditable AI methodologies
- integration of AI capabilities into regulatory monitoring and oversight.

Mining companies engaging proactively with regulators on appropriate AI applications may influence this evolution toward risk-based approaches that encourage innovation while maintaining rigorous safeguards.

AI models themselves will be maintained and improved. For example, as more data is collected (new inspection reports, new sensor data), the models can be retrained or fine-tuned to stay current. If engineers embrace the system they can correct or validate the AI’s outputs, those learnings can be used to enhance the model’s accuracy over time. In essence, creating a dynamic system that gets better and smarter as it’s used.

7 Conclusion

The application of artificial intelligence to tailings knowledge management systems represents a significant opportunity to enhance GISTM implementation without requiring prohibitive investment. The case studies presented demonstrate that organisations across the operational spectrum – from single-site operations to multinational corporations – can leverage these technologies to improve decision-making and strengthen governance.

Several key principles emerge from successful implementations.

1. Right-sized solutions that match organisational capacity and requirements yield better outcomes than one-size-fits-all approaches. The continuum from simple document processing to enterprise knowledge platforms offers multiple entry points based on specific needs and resources.

2. Domain expertise integration remains essential for effective implementation. AI tools augment rather than replace technical knowledge, with the most successful deployments combining tailings expertise, information management disciplines, and AI capabilities.

3. Incremental implementation delivers earlier benefits and reduces implementation risk. Organisations that began with focused applications addressing specific GISTM requirements achieved faster returns and built organisational confidence for expanded deployment.

4. Governance frameworks that clearly establish appropriate use cases, verification requirements, and decision authorities are essential for responsible deployment, particularly in high-consequence domains like tailings management.

These principles provide a foundation for organisations at any stage of their digitalisation journey to effectively leverage AI for enhanced tailings management. The democratisation of these technologies makes implementation increasingly feasible even for operations with limited resources, offering a practical path toward more robust knowledge management across the industry.

As GISTM implementation continues to advance globally, AI-enabled knowledge management represents not merely a technological convenience but an important capability for achieving the standard's fundamental objective – safer tailings facilities through comprehensive understanding, rigorous management, and transparent governance.

References

Anthropic 2023, *Claude 2*, large language model, https://www.anthropic.com/news/claude-2

Barnewold, L & Lottermoser, BG 2020, 'Identification of digital technologies and digitalisation trends in the mining industry', *International Journal of Mining Science and Technology*, vol. 30, no. 6, pp. 747–757.

Barros, DA, Pereira, JAA, Ferreira, MM, Silva, RA & Nascimento, GO 2020, 'Integrated management of mining enterprises to achieve the Sustainable Development Goals', *Sustainability*, vol. 12, no. 5.

Bommasani, R, Hudson, DA, Adeli, E, Altman, R, Arora, S, von Arx, S, … Liang, P 2022, 'On the opportunities and risks of foundation models', *Journal of Machine Learning Research*, vol. 24, no. 166, pp. 1–212.

Chalkidis, I & Kampas, D 2019, 'Deep learning in law: early adaptation and legal word embeddings trained on large corpora', *Artificial Intelligence and Law*, vol. 27, no. 2, pp. 171–198, https://doi.org/10.1007/s10506-018-9238-9

Chiesa, C, Robertson, A, Waygood, C & Tingey, S 2021, 'An industry guide to the Global Industry Standard on Tailings Management', *Mine Environment Management*, special issue on tailings management.

Chui, M, Manyika, J, Miremadi, M, Henke, N, Chung, R, Nel, P & Malhotra, S 2018, *Notes from the AI frontier: Insights from hundreds of use cases*, McKinsey Global Institute.

Davenport, TH & Ronanki, R 2018, 'Artificial intelligence for the real world', *Harvard Business Review*, vol. January–February 2018, pp. 108–117.

Deloitte 2022, *The AI Dossier: Managing Risk and Trust for the Top Uses for AI in Every Major Industry*, London, viewed 16 June 2025, https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-ai-dossier-full-report.pdf

Ghasemaghaei, M & Calic, G 2020, 'Assessing the impact of big data on firm innovation performance: Big data is not always better data', *Journal of Business Research*, vol. 108, pp. 147–162.

International Council on Mining and Metals, United Nations Environment Programme & Principles for Responsible Investment 2020, *Global Industry Standard on Tailings Management*, https://globaltailingsreview.org/wp-content/uploads/2020/08/global-industry-standard-on-tailings-management.pdf

Goebel, M & Fülöp, L 2020, 'The value of information management in the mining industry', *Mining, Metallurgy and Exploration*, vol. 37, pp. 1825–1835.

Jha, SK, Dwivedi, A, Tiwari, MK & Rana, NP 2022, 'AI-enabled knowledge management for sustainable mining: a research agenda', *Journal of Cleaner Production*, vol. 330.

Johnson, AE, Bulgarelli, L & Pollard, TJ 2021, 'Artificial intelligence in pharmaceutical research and development: a review', *Drug Discovery Today*, vol. 26, no. 1, pp. 80–88.

Katic, M & Majstorović, V 2022, 'Knowledge management in high-risk industries: a quantitative analysis', *Safety Science*, vol. 146.

McKinsey & Company 2023, *The State of AI in 2023: Generative AI's Breakout Year*, McKinsey Global Survey, https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year

Morgenstern, NR, Vick, SG, Viotti, CB & Watts, BD 2016, *Report on the Immediate Causes of the Failure of the Fundão Dam*, Fundão Tailings Dam Review Panel.

Mikalef, P, Pappas, IO, Krogstie, J & Giannakos, M 2018, 'Big data analytics capabilities: A systematic literature review and research agenda', *Information Systems and e-Business Management*, vol. 16, no. 3, pp. 547–578.

OpenAI 2023, *GPT-4*, large language model, https://openai.com/gpt-4

Prabhakar, A, Mosur, V, Cox, D 2021, *Rethinking your approach to AI*, IBM Institute for Business Value, Armonk, https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ai-business-strategy

Samimi, A & Najafzadeh, M 2021, 'Return on investment for digital technologies in mining: A quantitative analysis', *Resources Policy*, vol. 71.

Schoenberger, E 2016, 'Environmentally sustainable mining: the case of tailings storage facilities', *Resources Policy*, vol. 49, pp. 119–128.

Touvron, H, Lavril, T, Izacard, G, Martinet, X, Lachaux, M-A, Lacroix, T, ... Scialom, T 2023, *LLaMA 2: Open Foundation and Fine-Tuned Chat Models*, Large language model, Meta, https://ai.meta.com/llama/

Vergara, D, Rubio, MP & Lorenzo, M 2021, 'On the use of digital twins in mining operations', *Applied Sciences*, vol. 11, no. 14.

Vidal-Silva, C, Paredes-Belmar, G, Serón, J, Madariaga, E & Rodríguez, R 2021, 'Knowledge management systems and analytics in the mining industry: a systematic literature review', *Sustainability*, vol. 13, no. 4.