



White paper

Iron Mountain InSight[®]: AI model training, security and compliance



Executive summary

Iron Mountain InSight® empowers businesses to turn their information into actionable data through a secure, scalable platform. InSight is a low-code, software-as-a-service (SaaS) platform that ingests and processes unstructured content and can integrate with key business processes and systems via pre-built connectors or application programming interfaces (APIs). The platform processes DOCX, TIF, and PDF files across industries and functions, including HR, invoice processing, auto lending, healthcare, finance, insurance, and government. InSight is a portfolio and functionality may vary by product. Iron Mountain can work with you to support your organization's needs.

This whitepaper delves into the critical aspects of the AI-based functionality available in the InSight platform, related data privacy and security, as well as relevant customer's responsibility and Iron Mountain's liability. It provides a comprehensive overview of how Iron Mountain ensures responsible and secure AI deployment while adhering to evolving regulations (such as the EU AI Act) and best market practices.

I. AI-based functionality

AI models:

InSight leverages the power of large language models (LLMs) to enhance its capabilities, particularly in auto-classification and auto-extraction processes. The platform utilizes the "model attention" mechanism on specific document contexts for LLMs to effectively understand and extract key metadata fields. Additionally, InSight employs Retrieval Augmented Generation (RAG) to further enhance the LLMs' performance. RAG allows the system to retrieve relevant information from the customer's document repository, providing context-specific and accurate responses.

It's important to note that while the LLMs employed are broad in scope and not developed in-house by Iron Mountain, the company refines their application through prompt engineering and RAG to tailor the models to specific use cases. This approach combines the generative capabilities of LLMs with the ability to access and incorporate customer-specific information, resulting in more accurate and contextually relevant outputs.

Data input:

Iron Mountain utilizes a diverse range of data sources for AI model training, including the customer's document

repository, open-source data, internal data, and simulated data, based on the customer's scope of work. To ensure data quality and mitigate bias, the platform incorporates the following measures:

- **Human labeler evaluation:** Multi-labeler consistency methods ensure the accuracy of data labels.
- **Validation processes:** Techniques like mock data generators and Subject Matter Expert (SME) validation contribute to robust data validation.

Model training:

InSight offers a spectrum of AI model training solutions. Pre-trained AI models (both general and specialized) are available, along with model fine-tuning capabilities. Iron Mountain utilizes a blend of training methodologies depending on the specific model type and data availability.

Here's a breakdown of the training methodologies our implementation team employs:

- **Zero-shot/few-shot generative AI:** Leverages prompts and example context learning to guide the model towards generating desired results.

- **One/few shot CV solutions:** Ideal for scenarios involving a limited number of documents or templates, this approach utilizes document alignment and document fingerprint detection techniques.
- **Transformer/CNN based solutions:** Suitable for scenarios with a sufficient volume of labeled documents, this method involves fine-tuning deep neural networks.
- **Generative AI based solutions (small data):** Employs fine-tuning of large foundation models (large language and vision models) through techniques like prefix/prompt tuning and low-rank adaptation of large language models (LoRA).
- **Composite solutions:** Combines the aforementioned methodologies for optimal results.

Customization and user control

The platform's AI functionality can be turned on or off depending on the customer's specific requirements. Should the customer choose to turn off this functionality, they should contact their Iron Mountain account representative.

II. Data privacy and security

InSight prioritizes data privacy and security through robust measures:

- **Data privacy and security:** The processing of data is subject to the relevant data privacy and security standards following applicable regulations and best practices.
- **Encryption:** The platform employs industry-standard TLS 1.2 or higher and AES 256 encryption for data both in transit and at rest. This helps protect customer data from unauthorized access.
- **Exclusive deployment:** AI models are deployed within Iron Mountain's cloud infrastructure or on-premise*, guaranteeing exclusive access and preventing data exposure to third-party AI or cloud vendors.
- **Localized training:** InSight utilizes only the customer's own data for model training and validation, ensuring that one customer's data is never used to train models for another customer's.
- **Access control:** The platform implements multi-factor authentication and allows users to define data access boundaries, maintaining strict control over who can access sensitive information.
- **Data loss prevention:** The platform leverages data loss prevention (DLP) mechanisms to identify and mask personally identifiable information (PII) and preserve the original and de-identified copies, ensuring compliance with data privacy regulations.
- **Partner selection:** Iron Mountain partners with vendors that adhere to strict data security and privacy standards and controls (technical, administrative and contractual).

III. Ethical AI and bias mitigation

Iron Mountain employs a multi-pronged approach to promote ethical AI development and mitigate bias:

- **Diverse data sources:** The platform utilizes a range of data sources, including customer data, open-source data, and simulated data, to ensure model training is comprehensive and representative.
- **Human evaluation:** Multi-labeler consistency methods and subject matter expert (SME) validation are employed to ensure data quality and reduce bias in training datasets.
- **Continuous monitoring:** Automated systems and manual processes, including a dedicated SIEM system and CIRT team, continuously monitor for anomalies and potential biases.
- **Model prediction tracing:** Connecting model predictions to document resources to avoid unintended results.
- **Knowledge base and domain-specific rule engine:** Calibrating prediction confidence scores and guides interactions within the context of relevant documents.
- **Response filters:** Constraining interactions within the context of documents, preventing unintended outputs.

IV. Reliability and accuracy

Before production deployment, Iron Mountain rigorously tests model reliability and accuracy through:

- **Data collection and cleaning:** Gathering new and diverse datasets, including recent customer data, to ensure continued relevance and reduce noise or errors.
- **Rigorous testing:** New models undergo both quantitative and qualitative evaluation, including accuracy metrics, response latency testing, and human-in-the-loop (HITL) review.
 - » **Quantities evaluation:** Calculating accuracy metrics, response latency, and batch-data stress testing.
 - » **Qualitative evaluation:** Human labelers use human in the loop (HITL) to check prediction results and submit review reports.
 - » **A/B testing:** Comparing the performance of new models with existing models.
- **Model retraining:** Utilizing fine-tuning techniques to optimize model performance.
- **Thorough evaluation:** Rigorous assessment and comparison ensure the optimal solution is deployed.
- **Feedback loop:** User feedback is actively captured and incorporated to identify areas for improvement and enhance model performance. Users can submit tickets to report bugs or performance issues.
- **Industry standard accuracy metrics:** Utilizing established metrics to track model performance.

Error detection and correction

Prior to production deployment, Iron Mountain employs automated systems and manual processes to detect and correct errors or biases in AI models, such as those previously listed. These systems include:

- **Benchmark datasets:** Regular evaluation using benchmark datasets helps identify errors and biases.
- **Human-in-the-loop (HITL) correction:** Exception-based HITL flagging and correction mechanisms are used to address anomalies.

V. Integration and implementation

Iron Mountain offers flexible integration options:

- **Custom model development:** The platform supports custom model development through automated machine learning (AutoML) capabilities, allowing users to train models using their own datasets. The platform also provides a low-code labeling experience for creating datasets, making it easier for subject matter experts to train models without extensive technical expertise.
- **Deployment options:** InSight offers both cloud-based and on-premise deployment* options, catering to various organizational needs and compliance requirements.

VI. Transparency and explainability

Iron Mountain promotes transparency in its AI processes:

- **Model prediction tracing:** The platform connects model predictions to document resources.
- **Information sharing:** Users are provided with information about the model training and updating process, including details on training data methodologies used.

VII. Regulatory compliance

InSight is designed to meet various regulatory requirements:

- **Ongoing compliance:** The platform maintains compliance with evolving data privacy regulations like the General Data Protection Regulation (GDPR) and security standards such as ISO-27001 and SOC 2 Type 2.
- **Versatile support:** InSight successfully navigates diverse compliance landscapes, including GDPR in the EU, HIPAA and FedRAMP in the US.
- **Option to deploy on-premise*:** The platform offers on-premise deployment options for optical character recognition (OCR) and generative AI (genAI) solutions, catering to organizations with specific compliance requirements.

VIII. Customer's responsibility and Iron Mountain's liability

Customer's responsibility

- **Data provision and quality:** Customers are responsible for providing accurate, complete, and high-quality data for processing. The effectiveness of AI models, including AutoML and Generative AI, heavily depends on the quality of input data. Customers must ensure that their data is free from errors and inconsistencies to achieve optimal results.
- **Compliance with regulations:** Customers must ensure that their data processing activities comply with relevant legal and regulatory requirements, such as GDPR or California Consumer Privacy Act (CCPA). This may include obtaining necessary consents for data usage and ensuring data privacy.
- **Security measures:** While Iron Mountain provides a secure environment for data processing, customers are responsible for implementing their own security measures to protect data during transmission to and from the platform. This includes using encryption and secure access protocols.
- **Usage policies:** Customers should adhere to the platform's terms of service and usage policies. This involves using the AI tools provided by Iron Mountain ethically and responsibly, without engaging in activities that could harm other users or violate other's rights.

Iron Mountain's liability

- **Data security:** Iron Mountain is liable for maintaining a secure environment for storing and processing customer data. This includes implementing robust security measures to protect against unauthorized access, data breaches, and other cyber threats.

- **Model accuracy and performance:** While Iron Mountain employs advanced AI methods like AutoML & Gen AI as required pursuant to your SOW requirements and product specifications to enhance model accuracy, it cannot guarantee perfect results due to inherent uncertainties in AI predictions. Iron Mountain is committed to continuous improvement of its models to minimize errors.
- **System availability:** Iron Mountain ensures high availability and reliability of its platform services. In case of system downtimes or disruptions, Iron Mountain is responsible for promptly addressing issues to restore service continuity.
- **Support and maintenance:** Iron Mountain provides ongoing support and maintenance services to assist customers in effectively using the platform's features. This includes regular updates, troubleshooting assistance, and user training resources.

**On-premise deployment is available for select countries. For more information, please contact your local sales representative*

Summary

By leveraging robust security measures, transparent practices, and continuous improvement, InSight provides a secure and reliable platform for businesses to harness the power of AI while adhering to evolving regulatory requirements, industry standards and best market practices.



800.899.IRON | [ironmountain.com](https://www.ironmountain.com)

About Iron Mountain

Iron Mountain Incorporated (NYSE: IRM), founded in 1951, is the global leader for storage and information management services. Trusted by more than 225,000 organizations around the world, and with a real estate network of more than 98 million square feet across more than 1,400 facilities in over 60 countries, Iron Mountain stores and protects billions of information assets, including critical business information, highly sensitive data, and cultural and historical artifacts. Providing solutions that include secure storage, information management, digital transformation, secure destruction, as well as data centers, art storage and logistics, and cloud services, Iron Mountain helps organizations lower cost and risk, comply with regulations, recover from disaster, and enable a more digital way of working. Visit www.ironmountain.com for more information.

© 2025 Iron Mountain, Incorporated and/or its affiliates ("Iron Mountain"). All rights reserved. Information herein is proprietary and confidential to Iron Mountain and/or its licensors, does not represent or imply an invitation or offer, and may not be used for competitive analysis or building a competitive product or otherwise reproduced without Iron Mountain's written permission. Iron Mountain does not provide a commitment to any regional or future availability and does not represent an affiliation with or endorsement by any other party. Iron Mountain shall not be liable for any direct, indirect, consequential, punitive, special, or incidental damages arising out of the use or inability to use the information, which is subject to change, provided AS-IS with no representations or warranties with respect to the accuracy or completeness of the information provided or fitness for a particular purpose. "Iron Mountain" is a registered trademark of Iron Mountain in the United States and other countries, and Iron Mountain, the Iron Mountain logo, and combinations thereof, and other marks marked by ® or TM are trademarks of Iron Mountain. All other trademarks may be trademarks of their respective owners.

