ETHICAL AI, RIM AND THE ROLE OF ONTOLOGIES



By Stephen Clarke

INTRODUCTION

This article discusses how ethical AI governance can be supported by records and information management (RIM) processes and practices, particularly the use of relational ontologies. The role of RIM in ethical AI governance is, and will increasingly be, vital in ensuring AI content is compliant with legislation, and that it demonstrates the characteristics of reliability, integrity, authenticity and usability as set out in RIM standards such the *ISO 15489 Standard for Records Management.*[1]

This narrative is deliberately general and aspirational. It aims to provide an opportunity for further thinking and debate around the ethical use of AI, and the vital role that records and information management plays in its governance. To do this, it summarises content from some of my previous articles on the topic of AI, RIM and ontology, it also introduces some of my more recent thinking on the topic. It starts by covering what ethical AI is, and the role of RIM in achieving these aims. It then moves onto a summary of the role of ontologies in RIM, why they can be the ultimate underpinning design model for AI, and why they are essential in RIM as we move from traditional manual paper paradigm practice towards digital-native, automated modes of operation. It then concludes with some brief discussion of ontology-based RIM approaches for ethical AI.

WHAT IS ETHICAL AI?

As artificial intelligence (AI) technologies proliferate across public and private sectors, ethical governance has emerged as a central concern. Organisations and regulators are scrambling to develop responsible AI principles, accountability frameworks, and compliance protocols.[2]

For the purposes of this article, I define Ethical AI as the practice of designing, developing, and deploying artificial intelligence systems in a way that aligns with fundamental human values and rights. It's about ensuring that as AI becomes more integrated into our lives, it does so in a manner that is fair, transparent, accountable, and beneficial to society. At its core, Ethical AI seeks to prevent harm and promote well-being.

Ethical AI governance provides the framework to operationalize these principles. It is the comprehensive system of policies, processes, and structures that an organization puts in

place to ensure its AI activities are conducted responsibly. This governance extends beyond simple legal compliance to embed ethical considerations throughout the entire lifecycle of an AI system, from its initial conception to its eventual decommissioning.

THE KEY ASPECTS OF ETHICAL AI GOVERNANCE

The practice of Ethical AI and its governance revolves around several key principles:

- *Fairness and Non-Discrimination:* AI systems, particularly those used for decision-making in areas like hiring, lending, and criminal justice, must be free from biases that could lead to unfair outcomes for individuals or groups. This involves scrutinising the data used to train AI models for inherent biases and implementing techniques to mitigate them, and keeping full and accurate records of these activities.

- *Transparency and Explainability*: For an AI system to be trusted, its decision-making processes should not be a "black box." Stakeholders should be able to understand, to a reasonable degree, how an AI model arrives at its conclusions. This is crucial for identifying errors, challenging biased outcomes, and ensuring accountability. Robust recordkeeping of these processes is a critical underpinning to achieving these outcomes over time.

- *Accountability and Human Oversight*: Clear lines of responsibility must be established for the outcomes of AI systems. This means that there should always be a "human in the loop" who can intervene, correct, or override an AI's decision, especially in situations where the outputs impact on humans' rights and entitlements. Organisations must be accountable for the impact of their AI technologies and must maintain records for this purpose.

- *Privacy and Data Protection*: AI systems often require vast amounts of data to function effectively. Ethical AI demands that this data, particularly personal information, is collected, used, and stored in a way that respects individual privacy and complies with data protection regulations. Having this information available for future audit purposes can be challenging.

- *Reliability and Safety*: AI systems must be robust and secure, performing as intended without causing unintended harm. This includes protecting them from malicious attacks and ensuring they function safely in real-world environments

In practice, Ethical AI governance involves establishing internal AI ethics apparatus, creating clear AI usage policies and codes of conduct, conducting risk assessments to identify potential ethical harms, and fostering a culture of responsible

AI development and deployment within an organisation. These practices mirror the existing apparatus and practice for information governance and there is a clear opportunity to align and/or combine the functions within established organisational governance structures.

THE CRUCIAL INTERSECTION WITH RECORDS & INFORMATION MANAGEMENT

The principles of Ethical AI are deeply intertwined with the RIM discipline. As AI systems generate and consume vast quantities of information, the role of RIM professionals becomes paramount in ensuring ethical governance. The interface between Ethical AI and RIM manifests in several critical areas, which I will briefly summarise.

High-quality, well-governed data is the bedrock of trustworthy AI. RIM professionals play a vital role in establishing and maintaining robust data governance frameworks. This includes ensuring data quality, managing 'data lineage' (provenance), and implementing appropriate access controls. Without sound data governance, the risk of "garbage in, garbage out" in AI systems, leading to biased and unreliable outcomes, is significantly heightened.

AI systems create new types of records, from the prompts used to interact with generative AI to the automated decisions made by machine learning models. These records must be managed to ensure their authenticity, reliability, integrity, and usability. RIM principles provide the foundation for identifying what AI-generated information constitutes a record, determining how long it needs to be retained for business, legal, or accountability purposes, and ensuring it is stored in appropriate recordkeeping systems, not just within the AI platform itself.

A key aspect of Ethical AI governance is the ability to demonstrate how an AI system was developed and how it makes decisions. RIM professionals are essential in ensuring that the entire AI lifecycle is properly documented. This includes records of the initial business case, the data used for training, the design and testing of the model, and any subsequent updates or changes. This documentation is crucial for audits, investigations into algorithmic harm, and maintaining public trust.

AI presents unique challenges for information management, including the potential for creating misinformation, infringing on intellectual property rights, and perpetuating biases present in its training data. RIM professionals are at the forefront of developing policies and procedures for the responsible creation, use, and management of AI-generated content, helping to mitigate these risks.

In the age of AI, the role of records and information management extends beyond traditional archiving. It is a strategic function that underpins the very possibility of ethical and trustworthy artificial intelligence. By applying the established principles of their profession to the new frontier of AI, RIM professionals are critical to ensuring that these powerful technologies are harnessed for the benefit of all and ensure transparency, accountability, and the minimisation of bias.[3]

THE ROLE OF RIM IN ETHICAL AI GOVERNANCE

Many of the challenges currently being addressed by those debating the ethical use of AI are not new. In fact, RIM professionals have often confronted the same ethical questions now being rebranded through an AI lens: how to ensure data quality, maintain transparency, protect privacy, and uphold public trust.

Far from being left behind, I believe that RIM professionals hold a key role in aiding ethical AI governance, both in terms of managing the risks and realising the benefits of AI. This belief is fuelled by my suspicion that much of the current discussion around the ethical management of AI systems are based on the often under-employed and/or under-appreciated RIM skills and archival processes as summarised above. As a result, RIM-related tools and thinking may offer practical pathways to AI accountability, explainability, and compliance.[4]

Traditionally records management practice has been designed for creating, maintaining, and ensuring the accessibility of information within a paper paradigm. For example, EDRMS and document management systems were created to address the increasing scale of electronic documents, and to provide electronic access across space (you no longer had to have access to a filing room in a specific location, or get files delivered to your in-tray). There were also collaborative systems (such as SharePoint) that were adapted to manage information to a certain level and achieve 'document management' capabilities. However, these systems are often not yet able to fully manage the sheer scale or complexity of current information, and to achieve 'records management' level capability. AI tools can finally address this issue. For example, where collaboration tools provide document management but not records management level attributes and outcomes at an evidential level, AI can minimise human error and deficiencies in manual classification and manage metadata attribution, consistent labelling, version control, etc, particularly if underpinned by the automation of RIM tools and practices such as ontology-based information modelling.

With the rapid uptake of AI, much of how we 'see' and treat information is being disrupted, and this is also true of both the public and private sectors with the growing use of GenAI business tools. While the Aotearoa New Zealand government is developing guidance,[5] the day-to-day use of AI can still often result in practices that not only plagiarise existing content, but raise several issues around information access, use, and dissemination as summarised above in the key aspects of AI Governance. As a result, the need for robust ethical frameworks and governance mechanisms around AI (particularly when used in business systems) becomes mission critical for modern organisations.

This is where RIM practice and practitioners come into their own. While we have seen the rise of (often self-proclaimed) new role descriptions, such as "AI Ethicists", and "AI Governance Experts", these roles are often based on a lack of awareness that many required solutions already exist in RIM. For example, I consider that RIM practice can support ethical AI through appropriate and accountable practices around data and information stewardship, integrity, process management, disposal, privacy and access. RIM practices in these areas are not new, being based on legislation such as the Public Records Act 2005,[6] and the Privacy Act 2022,[7] and established theory and practice. For example, the original Lifecycle Model was conceived by Phillip Coolidge Brooks and Emmett J. Leahy in the United States National Archives in the 1940's,[8] before itself being superseded in the 1990's by the Records Continuum Model to accommodate the non-linear nature of electronic and digital records.[9]

THE ROLE OF ONTOLOGIES IN RIM

Before delving into its pivotal role in ethical AI, it is necessary to understand the role of the ontology in RIM professional practice. Ontology provides a structured and conceptual framework for organising information, which improve the consistency, quality, and discoverability of data and information by creating a shared semantic layer between the data and information and the user, providing the following benefits:

- *Improved Data Quality*. Ontologies establish clear definitions for concepts, ensuring that data across systems is consistently structured and semantically aligned. This directly contributes to higher data quality, as it eliminates ambiguity and misinterpretation.

- *Enhanced Interoperability*. By using a common conceptual framework, ontology supports seamless data integration across diverse business systems. This is particularly valuable in large

organisations where business-related information often resides in silos.

- *Efficient Data Governance*. Ontologies assist in data governance by clearly defining the roles, rules, and relationships of data elements, making it easier to manage access, compliance, and security.

- *Better Knowledge Discovery*. With an ontology in place, organisations can leverage advanced search and reasoning capabilities. Users can query data in more sophisticated ways, discovering relationships and insights that may not be apparent through traditional keyword-based search.

Semantics

The role of semantics is central to ontologies due to its ability to model the meaning of data and information by focusing on how concepts, terms, and relationships are defined and understood. This includes semantic principles such as:

- *Defining Terms and Concepts*. Semantics ensures that each term or concept within the ontology has a clear, agreed-upon meaning. This is critical for enabling shared understanding between different systems or people.

- *Relating Entities*. Semantics in ontology modelling defines how entities are related (e.g., part-whole, type-instance, cause-effect relationships). This allows systems to interpret the meaning behind data, rather than treating it as isolated facts.

- *Disambiguation*. Semantics helps to differentiate between homonyms (words with multiple meanings) and synonyms (different words with the same meaning), ensuring that the data model captures the correct meaning in context.

- *Formal Semantics*.: Ontologies use formal logic to define the relationships between entities, allowing systems to reason over data and derive new meanings or facts based on the modelled relationships.

Semantics gives meaning to the concepts and relationships in an ontology, allowing for effective communication between systems and enabling machines to interpret the data in a human-like way.

Semantics and Tautology

Semantics also help ontologies to avoid the risk of tautology, which is saying the same thing in different ways (for example "free gift"). This ensures that relationships and definitions are logically consistent and non-redundant, aiding:

- *Consistency in Definitions*. Ontologies must avoid tautological definitions (where terms are defined using themselves), ensuring clarity and avoiding circular reasoning in data relationships.

- *Logical Integrity*. While tautology in natural language is often seen as redundant in formal logic, it can ensure that certain statements are always true (e.g., a "dog is a dog"). In ontology design,

tautological rules help maintain the logical integrity of the system, ensuring that any entity is correctly classified without contradiction.

- *Minimising Redundancy*. A well-constructed ontology minimises tautological or redundant information to avoid inefficiencies and confusion, ensuring that each concept is defined uniquely and with purpose.

The avoidance or minimisation of tautology ensures logical consistency within the ontology, avoiding circular definitions and maintaining a robust framework where entities and relationships are logically sound and non-redundant.

Epistemological principles

Epistemology focuses on the theory of knowledge, particularly around how we can clarify what we think we know. Within the context of data and information modelling, epistemology more specifically relates to how knowledge is structured, acquired, and validated within the system at a semantic level. As a result, ontologies can apply epistemological principles to classify or categorise different facets of information in order to help both people and machines differentiate between "what we know" and "how we know it". For example, by:

- *Knowledge Representation*: Defining the relationships between entities and concepts to represent knowledge in a way that both humans and machines can understand.

- *Assertions and Inference*: Ontologies allow for logical reasoning, where new knowledge can be inferred from existing data. This can come from embedded epistemological principles that ensure that the data structures support the derivation of valid conclusions and decision-making.

- *Trustworthiness*: Epistemological principles can also ensure that the provenance of data (where it comes from and how it was derived) is clear, making it possible to evaluate the credibility of the information.

Ontologies use epistemological principles to ensure that data is not just stored, but that its meaning and sources are well-understood. These principles shape how knowledge is represented, validated, and organised within the ontology, ensuring that it reflects the way knowledge is structured in the real world.

ONTOLOGY AS THE ULTIMATE UNDERPINNING DESIGN MODEL FOR AI

Due to the considerations and benefits outlined above, ontologies not only provide a strong foundation for RIM, but also RIM provides a strong foundation for AI governance when it is used in relation to business systems. For example, an ontology-based approach to RIM can support AI business systems in returning results that are clearer, more structured, and with a more semantically rich understanding

of the records. This is particularly important when using a query-based approach to apply records management policies such as retention and disposal.

Employing ontological approaches can also enhance better information governance around the use of AI through broader system interoperability, information compliance, contextual accuracy, and scalability. For example, ontologies can help with:

- *Knowledge Representation*: Defining the relationships between entities, and concepts to represent knowledge in a way that both humans and machines can understand. This helps systems differentiate between "what we know" and "how we know it".

- *Assertions and Inference*: Ontologies allow for logical reasoning, where new knowledge can be inferred from existing data. This helps ensure that the data structures support the derivation of valid conclusions and decision-making.

- *Trustworthiness*: Ensuring that the provenance of data (where it comes from and how it was derived) is clear, making it possible to evaluate the credibility of the information.

Ontologies can also ensure that relationships and definitions are logically consistent and non-redundant. For example, they can help with:

- *Consistency in Definitions*. Ontologies avoid tautological definitions (where terms are defined using themselves), ensuring clarity and avoiding circular reasoning in data relationships.

- *Logical Integrity*. Ontologies aid formal logic where required, ensuring that certain statements are always true (e.g., a "dog is a dog"). In ontology design, inherent rules help maintain the logical integrity of the system, ensuring that any entity is correctly classified without contradiction.

- *Minimising Redundancy*. A well-constructed ontology minimises vagueness or redundancy to avoid inefficiencies and confusion, ensuring that each concept is defined uniquely and with purpose.

Overall, these aspects help to avoid the use of circular definitions, maintaining a more robust framework where entities and relationships have a greater chance of being logically sound and non-redundant, and helping to better represent the way knowledge is structured and represents the real world.

Ontologies also provide value when customising AI models like Generative Pre-Trained transformers (GPTs), Large Language Models (LLMs), and in particular business-centric Small language Models (SLMs) in making them more relevant and effective in specific cultural and organisational contexts. We currently see this challenge with different GPTs across cultures e.g.

Mistral from the EU, DeepSeek from China, or Claude from the USA. The underlying cultural perspectives, socio-political perspectives and societal mores increasingly present an emerging frontier for understanding information context across cultural realms.

Accordingly, for organisations to derive the most value from these models within the business environment, they need to be fine-tuned to specific contexts—both organisationally and culturally. Ontologies can help to create culturally specific frameworks that define key concepts, practices, and language use. They may also (eventually) allow generally accepted conceptual understandings to be mapped across culturally specific concepts. For example, an ontology tailored to a particular cultural context may be able to 'teach' an AI system the nuances of communication, norms, and values that are unique to that culture. For example, a science-based organisation with a highly specialised domain of knowledge a general-purpose language model will usually lack the precise understanding of the technical language and interdependencies within that domain. By incorporating the organisation's ontology, an AI system could better grasp these nuances within existing business systems, hopefully leading to more relevant outputs and more accurate decision support.

In addition, ontology-driven fine-tuning can help organisations embed their internal policies, values, and mission-critical objectives into AI models. This ensures that any AI-driven automation or decision-making processes are aligned with the organisation's strategic priorities, not only minimising risks and enhancing efficiency, but also enhancing AI systems to be more culturally aware and contextually appropriate. This could be particularly valuable in applications like customer support, content creation, and global-scale decision-making.

Knowledge graphs

In addition to the benefits derived from ontologies themselves, ontology-based knowledge graphs also play a critical role in supporting better AI logic processing using business systems. For instance, they can serve as a practical, flexible and dynamic option for representing the relationships and hierarchies defined in the ontology, enabling AI systems to interact with, infer from, and visualise the structured relationships between an organisation's records and its associated data and metadata. This is because knowledge graphs organise information into nodes (entities) and edges (relationships), making it easier for AI to traverse and query complex data structures.

Knowledge graphs can also, due to their dynamic, integrated

and extensible nature, allow for continuous updates and integration of new data. This enables the better support of evolving organisational needs, enhancing RIM maturity and capability and helping AI systems perform complex, context-aware queries. For example, identifying relationships between records or extracting trends across datasets. Four practical implications of ontologies include:

1. *Enhanced Personalisation*. Ontology enables LLMs to generate more relevant, contextually aligned outputs, whether the goal is customer engagement, internal communication, or strategic analysis.

2. *Improved Collaboration* Across Teams. Ontology can serve as a shared knowledge framework, helping cross-functional teams communicate more effectively and work together on AI initiatives with a common understanding.

3. *Ethical AI*. Ontology helps organisations develop AI systems that are more culturally sensitive, ethical, and aligned with their values. This is particularly important in avoiding biases and ensuring fair, inclusive AI applications.

4. *Faster Time-to-Value*. With an ontology in place, organisations can accelerate the fine-tuning of AI models. The time and effort needed to adapt a generic LLM to an organisation's needs are significantly reduced, allowing businesses to extract value from AI more quickly.

WHY RELATIONAL ONTOLOGIES ARE ESSENTIAL FOR RIM

Traditional hierarchical taxonomies have long served as a useful basis for data and information management tools in RIM (for example, by supporting function-based classification systems). However, they are increasingly proving insufficient to handle the complexity and dynamism of modern data and information environments, particularly in the private and public sectors. Key limitations hierarchical taxonomies are:

1. *Rigidity*: Taxonomies are typically static. Once data is categorised, it's difficult to adapt to new information, relationships, or changes in the way data is used or understood.

2. *Loss of Context*: Hierarchical taxonomies often force data into narrow categories, which can result in a loss of the rich, contextual relationships that exist between entities.

3. *Over-simplification*: Real-world data is rarely as clean and straightforward as a taxonomy suggests. Complex entities often belong to multiple categories or have relationships.

4. *Poor Adaptation to Change*: In a taxonomical structure, adding new data or understanding new relationships may require reworking entire hierarchies. For example, when new technologies or business processes emerge,

hierarchies can struggle to accommodate these shifts without undergoing significant restructuring.

5. *Inflexible for AI and Machine Learning*: AI systems thrive on understanding patterns, relationships, and context. Without a relational database to "translate", hierarchical taxonomies can limit these capabilities by constraining data into rigid categories, making it harder for AI to detect the subtleties and complex interconnections inherent in relational data.

6. *Fragmentation and Redundancy*: Hierarchical systems often lead to data fragmentation, where similar data is classified in different parts of the taxonomy without clear connections. This can lead to redundancy and inefficiencies, as information that could be linked together remains siloed.

As a solution, a wider, relational, ontological approach could offer further flexibility, interconnectedness, and adaptability in RIM-related tools and practices. It would also allow for the adoption of a more flexible, web-like structure where data points are related through multiple, often dynamic, associations. That is, entities are not forced into one-size-fits-all categories but instead understood in the context of its various relationships—whether those are temporal, geographical, functional, or thematic. Two key advantages of this approach are:

1. *Flexibility and Scalability*: A relational ontology can evolve with the organisation. As new data sources or relationships emerge, the ontology can easily adapt without the need to reclassify everything under a rigid hierarchy.

2. *Contextual Richness*: By capturing relationships between entities in multiple dimensions, a relational ontology offers a much richer representation of information, one that mirrors the complexity of real-world systems and interactions. This allows for more nuanced data discovery, analysis, and decision-making.

For organisations managing large volumes of data and information, a fully relational ontology approach could also help to:

- *Handle Complexity*: Organisations often deal with intricate and complex data ecosystems that include structured and unstructured data, multimedia assets, customer interactions, and Internet of Things (IoT) generated information. A relational ontology can manage these complex interconnections, enabling the discovery of insights that a taxonomical model might miss.

- *Enable Semantic Search and AI*: Relational ontologies are at the heart of semantic search engines and AI applications. Because they capture relationships between data points, they enable more intelligent querying, making it easier to locate

relevant information and generate insights. Ontology-driven AI systems can use this relational web to understand context, reason over data, and deliver more accurate predictions or recommendations.

- *Enhance Interoperability*: A fully relational ontology provides a shared conceptual framework that allows data from different systems to interact seamlessly. This is especially important in large organisations or multi-party collaborations where data silos can hamper efficiency. Relational ontology fosters interoperability by enabling data from various domains to connect in meaningful ways.

- *Allow Dynamic Data Governance*: A relational approach makes it easier to manage data governance in a world of ever-evolving privacy regulations, security requirements, and business needs. Instead of manually updating hierarchical categories and policies, organisations can automate governance processes based on the relationships between data elements. For example, if a new regulation affects data tied to a specific process, a relational ontology can instantly update governance protocols across all related data.

- *Tailor AI Systems to Organisational Contexts*: Relational ontologies enable AI systems to be fine-tuned to the unique processes, workflows, and culture of an organisation. By embedding specific relationships and knowledge structures, organisations can create AI models that are contextually aware and highly relevant.

- *Support Generative AI (GenAI)*: A relational ontology is crucial for enabling GenAI systems to generate contextually and culturally appropriate outputs. By embedding complex relationships and cultural nuances into the model, organisations can ensure that GenAI generates more accurate, meaningful content that aligns with their specific needs.

- *Improve Decision-Making*: With a relational ontology, decision-makers can query data from multiple perspectives, uncover hidden relationships, and make informed choices that reflect the full complexity of the situation.

ONTOLOGY-BASED RIM APPROACHES FOR ETHICAL AI

For the reasons above, I think it becoming increasingly critical for us to have an ontology-based approach to RIM, particularly when it comes to the development of tools and systems that help with the efficient running of AI systems, and for the operation of the AI tools themselves for use in RIM systems or general business applications. With this in mind, I have outlined nine reasons below that explain some of the benefits of doing so.

1. *Provides Semantic Interoperability*. Organisation-based AI systems often operate with diverse datasets and information sources from across a variety of business systems. An ontology can provide a shared vocabulary and standardised structure for data representation to enable more seamless integration across different RIM systems and domains. It can also ensure that records from these disparate sources are interpreted in a more consistent and meaningful way by AI, avoiding some potential misinterpretation by users or the development of information/data silos.

2. *Enhances Contextual Understanding*. Ontologies can also define the relationships between entities, attributes, and processes within records, offering rich contextual information for AI systems to read and analyse, and more accurately apply classifications, policy decisions and any additional contextual metadata. These processes in turn can improve the decision-making capabilities in AI, helping it to grasp the nuanced meanings and relationships inherent in the records and to retain records of these contextual relationships, as well as more trustworthy evidence of these decision-making processes.

3. *Improves Data Governance and Regulatory/Legislative Compliance*. RIM practice requires compliance with a variety of legal, regulatory, and organisational policies.[10] Ontologies can encode rules and classifications that align with these requirements, making it easier for AI to automatically enforce compliance routines and policies within and across business systems. For example, through access permissions, privacy protocols, and legal holds/retention and disposal actions. Not only is this valuable to ensure ongoing business compliance, it also helps to enable better information-related compliance in highly regulated industries such as healthcare, finance, and law, where regulatory adherence is critical.

4. *Aids Efficient Records Classification and Retrieval* Ontology-based approaches allow AI to classify and retrieve records more effectively by leveraging the hierarchical and relational structure of information and also provide the basis for automation of these processes. This ensures that relevant records are surfaced based on their meaning and context rather than just through keywords, enhancing search precision and efficiency for users, as well as access controls, de-duplication and version control.

5. *Provenance Tracking and Process Automation.* Ontologies can provide the apparatus for assigning definitions and workflows around the creation, use and disposal of business records and information, enabling AI to automate and monitor them. Ontology-based frameworks can also facilitate the tracking of provenance, ensuring

transparency and accountability in how records are used and modified by AI systems. This is critical for implementing records and information management practice in AI Business Systems as we move from twentieth century models of lifecycle management to modern data flows were the status of information objects can move through all the phases of the lifecycle, in any direction based on forms of Continuum thinking.[11] This non-linear information model means the ability to capture provenance in ever more complex muti-entity relationships makes it even more crucial to track and capture provenance metadata effectively and retain it over time as part of the record context.

6. Facilitates Explainable AI (XAI). Explainability is critical for AI systems, particularly in high-stakes applications. An ontology-based RIM approach makes it easier to trace and explain AI decisions by linking them back to well-structured, semantically rich data sources. This improves trust in AI systems and enables stakeholders to audit decisions effectively, and retain trustworthy, accurate authentic and usable records of these business processes.

Essentially, we cannot have accountability expressed through governance, audit, or support explainability without the ability to consult the records. That is, the records of inputs, outputs, algorithms, training data, training methods, the determinations and subsequent actions. The maintenance and availability of records, with their contextual attributes, provide information on the process of quality of the underlying data and the logic and reasoning of engine performance, and are therefore critical to understanding the final results of an AI query or process and detecting any potential bias or inaccuracy.

7. *Supports Scalability and Adaptability*. AI systems will increasingly need to adapt to new domains, data types, or regulatory environments. Ontologies can be extended or modified to incorporate new knowledge without disrupting existing structures. This can make the RIM framework scalable and future-proof, thereby allowing for diversification and growth in, and across, AI capabilities and data complexity, extensibility across geographical spaces, languages and cultural frameworks at a more global scale.

8. *Addresses Ethical Concerns.* Ontologies can be used to define ethical guidelines and constraints directly within the RIM framework and associated information architecture, guiding AI systems to operate within pre-determined and well-scoped boundaries, helping to set and enforce boundaries and guardrails.

9. *Aids Enterprise Business Information Standardisation.* By adopting an ontology-based approach, organisations can create a standardised foundation for records management. This can help to reduce inconsistencies and ensure a more uniform application of AI processes across teams, departments, and regions. It would also provide a basis for better systems-based interoperability across organisations and more consistent metadata schema and data classification for cross-organisational information systems and data-lakes.

This approach would also allow for a centralised data dictionary, 'business taxonomy', or term-store. That is, if a new item is added or a term removed or altered, this change can then be applied globally across the data ecosystem - not just locally within the particular silo, single instance, or subset of the overall data enterprise information and data ecosystem.

CONCLUSIONS

In essence, ethical AI governance requires embedded RIM practice, and RIM practice enabled through a relational ontological framework. In combination, these approaches enable organisations who use (or plan to use) AI to develop forms of good and sustainable data and information practice that by design incorporate concepts of data quality, privacy, compliance, and transparency and accountability.

As AI continues to evolve, integrating ontological concepts around established records and information management principles are vital in order to develop workable, responsible, effective and ethical AI solutions through traditional RIM capabilities, particularly since the use of AI technologies in business are only going to expand. RIM concepts, models and approaches are sound but need to be implemented in ways that support both AI automation of RIM practice, and RIM Practice for AI systems and outputs.

ENDNOTES

1. Standards Australia, "Information and documentation – Records management. Part 1: Concepts and principles (AS ISO 15489.1: 2017)", 2017.
2. See for example, Archives New Zealand, 'Atamai hangahanga me ngā mauhanga tūmatanui - Artificial intelligence and public and local authority records', 2023, available at https://www.archives.govt.nz/manage-information/how-to-manage-your-information/implementation/artificial-intelligence-and-public-records, accessed 3 July 2025.
3. For example, Floridi, L. (2010–2022) wrote various works on the ethics of information and the philosophy of the infosphere. See for example, L Floridi, The Ethics of Information, Oxford University Press, 2013 https://doi.org/10.1093/acprof:oso/9780199641321.001.0001. See also L. Floridi, The Logic of Information: A Theory of Philosophy as Conceptual Design. Oxford University Press, 2019. https://doi.org/10.1093/oso/9780198833635.001.0001.
4. L Floridi and J Cowls, 'A unified framework of five principles for AI in society', Harvard Data Science Review, vol 1, no.1, 2019. https://doi.org/10.1162/99608f92.8cd550d1
5. See for example Digital.Govt.NZ. n.d. "Artificial Intelligence (AI)." n.d. https://www.digital.govt.nz/standards-and-guidance/technology-and-architecture/artificial-intelligence.
6. Public Records Act 2005, available at https://www.legislation.govt.nz/act/public/2005/0040/latest/DLM345529.html, accessed 8 July 2025.
7. Privacy Act 2020, available at https://www.legislation.govt.nz/act/public/2020/0031/latest/LMS23223.html, accessed 15 July 2025.
8. See for example P C Brooks, 'What Records Shall We Preserve?', 1940, available at https://babel.hathitrust.org/cgi/pt?id=pur1.32754079146159&seq=3. Here he introduced the concept of the 'life history' of records, a concept later generally called the 'life cycle of records'. See also P.C Brooks, 'The Selection of Records for Preservation', American Archivist, vol. 3, no. 4, October 1940, pp. 221– 224, available at https://doi.org/10.17723/aarc.3.4.u77415458gu22n65.
9. See for example, C. Hurley, S. McKemmish, B. Reed, and N. Timbery, 'The Power of Provenance in the Records Continuum', Archival Science, vol. 24, no. 4, December 2024, pp. 825–45, available at https://doi.org/10.1007/s10502-024-09463-9.
10. See for example, J M Schwartz and T Cook, 'Archives, records, and power: The making of modern memory', Archival Science, 2002, vol. 2, no. 1–2, pp.1–19.
11. For background on the Continuum model that is the basis of this thinking, see Wikipedia, "Records continuum model", 16 October 2024, available at https://en.wikipedia.org/w/index.php?title=Records_continuum_model&oldid=1251586687