Title	RP04 FINAL Report: Preserving AI Techniques as Paradata
Working group code	WG2
Study title	Preserving AI Techniques as Paradata
Status	Final
Version	2
Writers	Scott Cameron, Patricia C Franks (lead), Kaila Fewster, Isto Huvila, Norman Mooradian
Use Case Contributors	RP04 Members: Mario Beauchamp, Souvick Ghosh, Rae Lynn Haliday, Nancy Powell, Alex Richmond CU05 Members: Maria Guercio, Stefano Allegrezza, Francesca Magnoni, Grandi Massimiliano, Bruna La Sorda, Maria Mata Caravaca
Date	March 30, 2026

Preserving AI Techniques as Paradata

RP04- Final Report

An ITrust[AI] Report
March 30, 2026

Preserving AI Techniques as Paradata

Executive Summary

Introduction/Problem Statement

Preserving AI Techniques as Paradata is one of more than 40 studies conducted by members of InterPARES Trust^{AI}—a multi-national interdisciplinary research project to design, develop, and leverage Artificial Intelligence to support the ongoing availability and accessibility of trustworthy public records.

Members of the Paradata study first met in 2021 to address a central question: When AI is used to support or automate archival, recordkeeping, or other processes, how much of the AI system should be preserved? Specifically, this might include the AI method, its code, training data (often drawn from existing records), test cases and results, parameters and their values over time, and the technical environment in which it operates.

Based on early examination of the AI process, the team quickly realized that records, metadata, and explainable AI (XAI) were not sufficient. Nor was a singular focus on the technical aspects of the AI process. The team determined that operational aspects of the AI process were equally important and that the missing element when documenting the AI process in order to facilitate transparency and accountability is Paradata. The team defined "Paradata" as "the information about the procedure(s) and tools used to create and process information resources, along with information about the persons carrying out those procedures" (ITrustAI glossary, https://interparestrustai.org/terminology/term/paradata).

Methodology

Researchers used a qualitative, iterative approach to their work. Among the data sources were one survey, two literature reviews, and four use cases. Because the term paradata was new in relation to the AI process, the researchers continuously shared their findings with a larger audience through peer-reviewed journal articles, presentations, professional articles, blog posts, and this final report.

Major Findings

Throughout its course, the study moved from a broad, conceptual analysis in its early stages, focusing on existing AI model documentation approaches, towards perspectives emphasizing contextually-situated assessments of paradata requirements designed for specific use cases which document the responsibilities and actions of individuals, systems, and organizations throughout the AI lifecycle. The following key points were identified:

1. Early conceptual work revealed that paradata overlapped significantly with the broad category of metadata, but that most AI transparency requirements would require greater volumes of information to be recorded than possible via most standard structured metadata schemata. As such, a firm boundary between paradata and metadata was not established.

2. Throughout the study, researchers emphasized the relevance of AI governance and legislation for assessing paradata requirements. Since most AI governance is rooted in a risk management framework, paradata can be expected to emerge throughout the AI lifecycle as an expected by-product of active and ongoing risk management strategies.

3. Since risk management strategies necessarily include documentation requirements, a survey of relevant AI governance documents uncovered a common set of facets across which necessary paradata may be evaluated. Paradata may be explicitly or implicitly required to satisfy accountability requirements; it may be generated by humans or automatically by AI tools; and it may be generated intentionally to directly convey information or as an incidental by-product of AI processes. Common documents arising from the AI implementation process—including risk assessments, impact analyses, ongoing quality assurance documentation, and more—are analyzed in relation to these facets.

4. Four case studies are summarized in the report: automated contact extraction for digitized records at the St. Louis Zoo; LLMs for predicting oil price movements at the Bank of Canada; AI use across clinical healthcare settings (meta-analysis); AI for reconstituting digital records aggregations (InterPARES companion study).

5. Researchers reflected on the evolution of their thoughts and practices regarding paradata following the completion of the study, and a summary of these reflections is included in this report. A key outcome is a transition in researchers' understanding of paradata as a static (but necessary) information object towards a process of document creation and preservation which must be built into AI implementations. Researchers appreciated the need for AI literacy on the part of archivists, who must be able to embed records capture processes early in the AI design process rather than after implementation.

Recommendations/Conclusions

As a takeaway for practitioners, the report offers a paradata reference chart which summarizes key points of any AI implementation which must be clearly documented; this document is offered with the caveat that specific applications will always possess unique paradata requirements.

The report also offers a conclusion and recommendations for further study, emphasizing the need for ongoing, fine-grained case studies within specific industries to build robust and detailed frameworks for broader applications.

Table of Contents

Introduction: Study Background

Preserving AI Techniques as Paradata is one of several studies conducted by members of the RP04 Working Group (WG), Retention and Preservation, led by WG Chair, Dr. Hrvoje Stančić. In September 2021, the InterPARES Trust AI Executive Committee approved the following abstract:

> If an AI technique is used to facilitate or automate an archival, recordkeeping, or other process, how much of that AI technique, its code, the data (probably a subset of existing records) we use to train it, test cases and test results to examine its efficacy, its parameters and their values at or over the time of application, the technical environment in which it is executive, and the records it (the AI technique) is applied to for automation purposes, should be preserved? This question is not about preserving AI techniques for their own sake, but about preserving them as contextual materials to support the protection of the records to which they are applied. As such, are they preserved as part of procedural context, technological context, a combination of the two, or other contexts? How do we preserve the pieces that constitute the AI technique (code, training data, test cases, parameters, etc.)? How reproducible should what we preserve be? If there is non-determinism or randomness in any of these AI techniques, how do we identify, characterize, and preserve them? If there is/are human(s) intertwined with the AI technique in the decision-making process, how is the human's role and his/her relationship with the AI technique captured and preserved? This study will explore these questions, gather data on the current state of practice, and propose best practices and solutions.

Drs. Patricia C. Franks of San Jose State University and Babak Hamidzadeh of the University of Maryland were appointed lead researchers for this study. The initial timeline was October 2021 through September 2024; however, the study was extended through the end of the InterPARES Trust AI grant period in 2026.

Developing a research question and a 3-year plan

The initial team comprised nine researchers, three consultants, and one graduate assistant, who met monthly; the team's size fluctuated over time due to the schedules of the individuals involved. At our first meeting, we refined a research question, as shown in the following paragraph, broadening it to encompass research on paradata for AI use, applied not only to archive and recordkeeping processes but also to other areas.

> If an AI technique is used to facilitate or automate an archive, recordkeeping, or other process, how much of that AI technique, its code, the data (probably a subset of existing records) we use to train it, test cases and test results to examine its efficacy, its parameters and their values at or over the time of application, the technical

environment in which it is executed, and the records it (the AI technique) is applied to for automation purposes, should be preserved?

The following three-year plan was devised to conduct our research:

2021-2022

The goal of this initial data gathering phase was to understand the current state of documentation of the AI process, particularly to identify paradata created and preserved in relation to archives and records management. We sought answers to questions like these:

- Are institutions using AI in their recordkeeping and/or archival processes?
- Do they document Paradata, such as AI code, training data, and case parameters?
- If so, how?
- If not, do they feel they should, and why?

To combine resources, the paradata team inserted questions into a survey being conducted by another research team, RP01 - Archival Challenges to be improved by AI. Respondents from the archives and records management fields were asked several questions recommended by members of our study.

The first: *Does your institution/organization use any automated or AI-supported activities in the digital preservation process?* Of the 106 respondents, 81% (86) replied no, 15% (16) yes, and 4% (4) don't know. Among the tasks cited by those who responded yes were ingest/upload/capture/packaging (3); classification (and granting access based on that) (2); metadata operations (description/extraction from records) (2); search and recommendations (2); and AI tools built into software applications (2). The second: *Could the AI-related technologies be integrated into the digital preservation system you are using?* Of the 106 respondents, 54% (57) said yes, 22% (23) said no, and 25% (26) did not respond. Figure 1 shows the activities identified by respondents as potential benefitting from the use of AI to improve the archival/records management process.

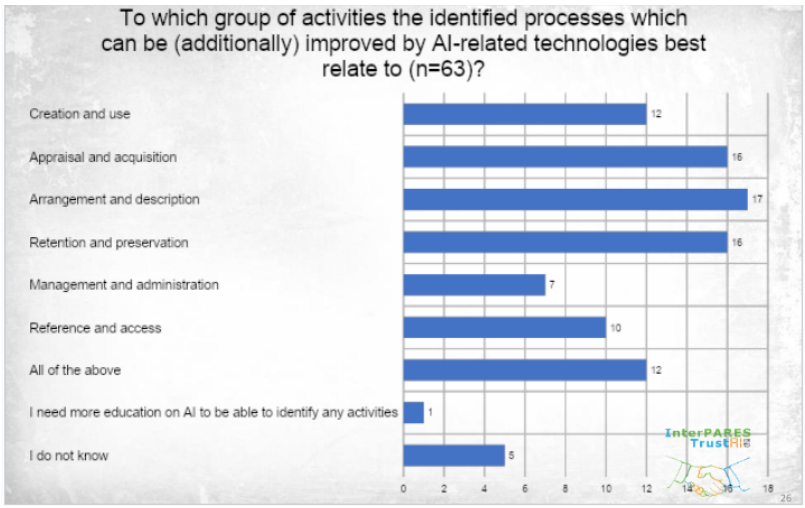

Figure 1. Responses to the survey question inquiring about archival / records management tasks that could be improved by the implementation of AI. *Source*: Presentation by Hrvoje Stančić and Željko Trbušić at the InterPARES Trust 5th Plenary Meeting, June 27-29,2022, Paris, France.

The results of the questions above indicate that archival and records management professionals were only beginning to use or explore the application of AI technologies and techniques in their daily practices. The most relevant question from this survey to our study is the following:

What information and documentation do you keep about your digital preservation processes? Describe or provide examples.

Because the term paradata was new and not yet understood in the context of AI, the word "documentation" was employed as a suitable replacement. Seventy-four of 105 respondents answered this question with mixed results. Ten reported having no information or documentation regarding their digital preservation process. Among the responses that might be considered metadata were 1) Database marking date, time, latitude, longitude, temperature, and preservation status of the document, and 2) Each document has its ID, date of creation, what it was created form, format, initial classification, notes and possible alterations, everything that happens with the document we can map and save. Sixteen of the 74 indicated that they used AI in the workplace, and only 13 of those responded to question 5.

The responses of only two of the 13 who indicated they used AI were consistent with the concept of paradata; see Table 1.

Table 1. Responses of two survey participants who used AI within their organizations and documented some of the process.

AI Employed and Reason	Documentation Created and Stored
Machine learning tool is used to assist archivists in identifying born-digital moving image records for preservation.	Manuals, process diagrams, technical specifications, etc.
Pattern recognition ML is used in the analysis of metadata for PII, and automated extraction of metadata from records using trained ML.	Extensive procedural documentation of all processes, related practices on the management of files created during digitization, system use and administration manuals, and IT security documentation.

Two lessons were learned as a result of this survey: 1) AI was not widely used by archivists and records managers, and 2) the documentation process, if any, was ad hoc. These findings indicated a need to disseminate information about both AI and the documentation of the AI process in the form of Paradata.

2022-2023

The goal of this second phase was to research potential practices and solutions by seeking answers to questions such as these:

- Should AI techniques employed be preserved as part of the procedural context, technological context, a combination of the two, or other contexts?
- How do we preserve the pieces that constitute the AI technique (code, training data, test cases, parameters, etc.)?
- How reproducible should what we preserve be?
- How do we identify, characterize, and preserve the AI technique if there is non-determinism or randomness in any of them?
- How do we document the human's role and relationship with the AI technique if one exists?

Our intent was to share our findings with the archives/records community and to formulate an approach to the preservation of AI techniques we could explore during 2023-2024.

2023-2024

During the third phase, the team intended to develop a pilot project to identify, capture, and preserve AI techniques as part of the preservation process. The goal was to provide a description and evaluation of the pilot project, share the results, and offer recommendations for further research. As will be explained in greater detail later in this report, our work plan was revised due to the task's complexity, and the study was extended beyond 3 years.

Defining the term Paradata when applied to the AI Process

After devising the plan of work, our next step was to define the term “paradata” as used in the study. Many definitions of Paradata were reviewed. The one below, drawn from the London Charter on the computer-based visualization of cultural heritage, served as the basis for our definition.

> Information about human processes of understanding and interpretation of data objects. Examples of paradata include descriptions stored within a structured dataset of how evidence was used to interpret an artefact, or a comment on methodological premises within a research publication. It is closely related, but somewhat different in emphasis, to "contextual metadata," which tends to communicate interpretations of an artefact or collection, rather than the process through which one or more artefacts were processed or interpreted. (https://londoncharter.org/glossary.html)

By the end of 2021, the study team agreed upon the following definition for Paradata in relation to documenting the AI process, which the Executive Committee approved for inclusion in the InterPARES terminology database:

> **Paradata** (n.) Information about the procedure(s) and tools used to create and process information resources, along with information about the persons carrying out those procedures (ITrustAI terminology, n.d.)
> https://interparestrustai.org/terminology/term/paradata

Initial Data Gathering Phase

Initially, the team formed three groups to address these tasks: 1) surveying the use of AI tools and technologies applied to archives and records management tasks—both from the vendor and the user perspectives, 2) reviewing at least one archives or records management process to determine where AI might be used to improve efficiency and efficacy and determine what paradata would be needed to document the process, and 3) performing use cases involving the use of paradata to document the AI process. The team later decided to focus on the first task, surveying the landscape, through the end of the 2021-2022 year.

Literature Review

An initial literature review was incorporated into a publication, “Positioning Paradata: A Conceptual Framework for AI Processual Documentation in Archives and Recordkeeping Contexts.” A final literature review was published in “Tracing the Past, Predicting the Future: A Systematic Review of AI in Archival Science.” The details of the literature reviews are provided in the corresponding documents. A summary of those foundational papers is included in this section.

Early literature - 2021-2023

In examining the literature at the start of our study, we quickly learned that the term paradata was used across several disciplines but not within the archival field. We also realized that the term was not used in the documentation of the AI process. Our initial findings were published in the paper "Positioning Paradata: A Conceptual Framework for AI Processual Documentation in Archives and Recordkeeping Contexts" (Cameron, Franks, and Hamidzadeh, 2023). The authors examined the use of paradata from three unique contexts: 1) as a term used in the social sciences, visual heritage, and documentation fields to describe information that records the process of creation or curation of other data or datasets. 2) in relation to existing XAI literature detailing some approaches to the Blackbox problem of understanding the logical pathways leading to an AI tool's decision, and 3) the placement of the term paradata and XAI in the literature on accountability and transparency in archives. The authors concluded there is a need for further research on AI documentation practices.

Several subsequent articles exploring the use of paradata to document the AI Process were published as part of our ongoing research. See *Appendix B, Dissemination*, for a detailed list. The fact that the first four articles returned in a recent Google Scholar search for the terms "Paradata + AI" were written by InterPARES Trust AI researchers supports the notion that we were exploring a new concept. A search for "Paradata + AI + archives produced the same results. A search for the terms "Archives + Paradata" without "AI" again returned the four articles written by the InterPARES Trust AI team, this time in positions 1, 2, 5, and 9.

Recent Literature - 2024-2025

Shinde, Kirstein, Ghosh, and Franks conducted further research in 2024 to examine how the topic "documentation of AI practices" had evolved since the study began in 2021. The results were published in a paper, "Tracing the Past, Predicting the Future: A Systematic Review of AI in Archival Science, delivered at the 2025 ASIS&T Annual Meeting and published in the Proceedings of the Association for Information Science and Technology (Shinde et al. 2025). A unique approach to gathering information for this paper was to identify five primary AI techniques used to enhance archives and records management practices, further divided into seven subthemes. The five themes and their subthemes are:

1. Natural Language Processing: Named Entity Recognition, Handwritten Text Recognition, Named Entity Disambiguation, and Information Retrieval
2. Traditional Machine Learning: Machine Learning
3. Deep Learning: Deep Neural Networks and Task Optimization
4. Computer Vision: Object Detection and Classification
5. Explainable AI (XAI): No subtheme

The next step was to categorize the most pertinent research on the use of AI to accomplish AI-enabled Records Management tasks. The four main themes with two to four subthemes each are as follows:

1. Records Management - Classification, Data Collection & Management, Retention and Disposition, and General Records Management.
2. Archival Processing - Appraisal, Arrangement and Description, and Preservation.
3. Access and Use - Privacy, Records Retrieval, Use of Records in Research, and Public Access
4. Professional Perspectives - Archival Education, User Perceptions.

As a result of this literature review, we found that research to that point in time, except for work coming out of the InterPARES Trust AI research project, could be divided into two main categories: 1) how archives work, and 2) the theory of AI tools and techniques. This situation clearly called for collaboration between archivists and computer scientists/AI developers to explore ways to improve archival practices using advanced AI techniques, an approach pursued by InterPARES Trust AI members.

Study Team

The following individuals contributed to the Preserving AI Techniques as Paradata research study.

Researchers	
Pat Franks (Study Lead)	
Mario Beauchamp	Babak Hamidzadeh
Scott Cameron	Isto Huvila
Fred Cohen	Norman Mooradian
Vic Ghosh	Alex Richmond
Rae Lynn Haliday	Corinne Rogers
Graduate Student Assistants	
Jeremy Davet (UBC GAA)	Kat Hodgson (UBC GAA
Kaila Fewster (UBC GAA)	Nancy Powell (SJSU GRA)

Part 1 - Paradata for the AI Process, the Concept - purpose of paradata

Paradata is a term used across multiple social science research contexts to describe diverse types of information that capture and document research processes. In archival and records contexts, our study has used the term to describe information recorded to elucidate complex computational processes that have influenced records or fonds. In this sense, we draw directly from the term's use in the statistical social sciences and in virtual heritage visualization. Across

contexts, paradata refers to information used to clarify and communicate the processes that have influenced other information resources (Cameron et al., 2023, 75:4). Our study is more narrowly focused on paradata in relation to AI in archives. From this pragmatic perspective, paradata is most relevant in high-risk AI applications in which the computational techniques employed are not fully transparent, understandable, or explainable. We posit that paradata can be introduced to highlight these gaps in understanding and to pose the question of what information might make AI processes understandable.

In the information field, paradata inevitably overlaps with the broader category of metadata. Defined informally and broadly as "data about data," information professionals frequently use metadata to describe structured, descriptive, and curatorial information recorded in a standardized, typically machine-readable structure. Much as metadata is a variable term with broad and narrow definitions, paradata may be approached with a narrow or a capacious scope. If we define paradata as information about the processes applied to other information resources throughout their lifespan, we are dealing with an expansive category that may be difficult to clearly delimit. In practice, paradata will likely be more narrowly delineated by the processual information required for a given context and will overlap significantly with structured metadata. As with descriptive metadata, the problem of capturing useful paradata often requires reducing the available information to a manageable volume. Accordingly, the boundaries between paradata and metadata are not easily delineated.

At the conceptual level, paradata is a type of information about other information resources—and hence a type of metadata, as a narrower set of "data about data" which serves the specific goal of enabling insight into otherwise obscure processes. Defining metadata more narrowly as structured, machine-readable information about other resources, paradata may appear both as elements within the structured metadata (for instance, within curatorial metadata such as PREMIS) and as other records or records groups in their own right, which may be understood as paradata in relation to the original resource. From this capacious perspective, paradata consists of information about another information resource which elucidates processes related to those information artefacts. This capacious definition of paradata is broader than the narrow sense of metadata as structured and encoded data about data and narrower than the broadest sense of metadata as data about data. Since paradata is so closely defined by its purpose (i.e., that of transparency into processes), the scope of what comprises paradata and what does not is dependent entirely on a given application process. As such, while a high level conceptual analysis of paradata can help scope the problem at hand regarding AI applications in archives, only a close analysis of AI use cases in context can give a clear understanding of its relevance and value to archivists and researchers.

Part 2 - Development of a consolidated framework for AI Risk Management

A key aspect of this study was to investigate how paradata could be integrated into the development of consolidated frameworks for AI risk management, particularly when these

tools are used for archival and recordkeeping functions, as a meaningful response to the fragmentation of current AI governance approaches. In many organizations, we found technical risk controls, legal compliance functions, and enterprise risk management structures for AI processes often operate in relative isolation, if they exist at all. The introduction of a consolidated framework seeks to integrate these domains into a coherent, lifecycle-based system grounded in traceability, accountability, and continuous oversight to more closely monitor AI processes and also align with the AI lifecycle. Furthermore, this type of integrated risk management structure aligns with emerging risk-based AI governance models, such as the NIST AI Risk Management Framework and the EU AI Act, which emphasize documentation, monitoring, and demonstrable control. In this sense, a consolidated AI risk management framework comprises three mutually reinforcing layers:

1. Governance and institutional oversight;
2. Technical documentation and paradata integration; and
3. Ongoing monitoring of AI tool development and use.

At the governance layer, responsibility for AI systems and their usage must be clearly allocated and defined. From AI model developers to risk and compliance officers, and ultimately users, each party involved in the development and deployment of AI tools should have defined roles across approval, review, and incident response processes. Moreover, these responsibilities should be embedded within existing enterprise risk management structures to prevent AI oversight from becoming siloed from other documentation processes within the organization. Standardized documentation requirements of organizational metadata should include the model's purpose, design plans, technical specifications, known model limitations, risk assessments, and procedural documents outlining how the model is intended to be used. Furthermore, formal checkpoints at key lifecycle stages (pre-deployment approval, material modification review, periodic reassessment) are important for improving accountability and reducing ambiguity in decision-making authority.

The technical documentation and paradata integration layer functions as the evidentiary infrastructure of the framework. Systematic capture of technical paradata, such as the code used to develop the models as well as training datasets, testing and validation results, version histories, performance logs, and evaluation and performance metrics help to create a continuous record of the system's behavior and context of use. By embedding automated logging and paradata-capturing mechanisms within model development and deployment pipelines, organizations can shift from retrospective reconstruction of risk events to real-time observability. In this sense, paradata supports verification: it enables organizations and, where relevant, regulators to investigate and assess the trustworthiness and robustness of a model and its corresponding levels of human oversight.

Finally, the ongoing monitoring layer operationalizes risk detection and response. Continuous performance evaluation enables early identification of model inconsistencies or degradation. Bias and fairness metrics can be monitored against predefined thresholds, triggering structured review mechanisms when deviations occur. Incident reporting mechanisms, ideally linked to

overarching governance bodies, ensure that identified risks are neither suppressed nor informally resolved.

A critical feature of a consolidated risk management framework is lifecycle continuity. Pre-deployment organizational paradata and deployment-phase technical paradata can be interconnected rather than compartmentalized to ensure continuous accountability across an AI model's lifecycle and prevent governance gaps during system updates or model repurposing. Risk, in this sense, therefore, is not confined to the model's initial design phase, but instead evolves alongside changing applications, user behaviors, and institutional contexts. Although operationalizing this type of risk management framework requires sustained institutional commitment, maintaining traceable paradata records across these transitions preserves organizational memory and strengthens both internal oversight and external accountability.

Part 3 - Classification and Mapping of Paradata to the AI process

The working group identified numerous artifacts that may constitute paradata. They also identified facets of paradata. A team analyzing authoritative documents to understand their explicit and implicit records requirements and the role of paradata in support of such records identified three facets that characterize paradata (Cameron et al., 2025).

To understand how paradata arises in the development and deployment of AI systems, the working group examined the AI lifecycle in many of its projects. Different lifecycle models were reviewed. The team found that paradata is part of the entire lifecycle, though different types and facets are more or less strongly represented across stages. Below is a description of different facets of paradata:

Explicit / Implicative Terms

Cameron et al. noted that authoritative texts contain terms that explicitly refer to information artifacts that can be considered paradata, as well as terms that imply information artifacts without explicitly identifying them. Examples of explicit terms are documentation, record, assessment, and plan. Examples of implicative terms are testing, training, and auditing. These terms are typically action terms that imply the creation of a document or record. The facet representing these categories in the table below is "Interpretation."

Human-Made / Automated

Another important classification in paradata is that between human-made information artifacts and machine-generated data and information. Cameron et al. found examples of both in their study of authoritative documents. Examples of human-made paradata include plans and

training manuals. Examples of machine-generated paradata include audit trails and output from testing programs. The label for this facet in the table is "Generation."

Intentional / Incidental

A third distinction in paradata identified by Cameron et al. is between paradata that is intentionally created to establish a record or to create context for a record and paradata created for a different purpose within the project but that can serve as a record or context for a record. The former they denominate as “Intentional”, the latter “Incidental.” The facet label for this distinction is “Purpose.” Unlike the other facets, this one is non-binary.

AI Development Lifecycle: CRISP-DM

Throughout the project, the team referenced various AI models. For the purposes of this mapping, we use the Cross-Industry Standard Process for Data Mining (CRISP-DM) model, see Figure 2. It is an established model used in data science and AI projects and the foundation of other AI lifecycle models. Table 2 maps paradata artifacts and their facets to the CRISP-DM model. Similar mappings can be made for other models.

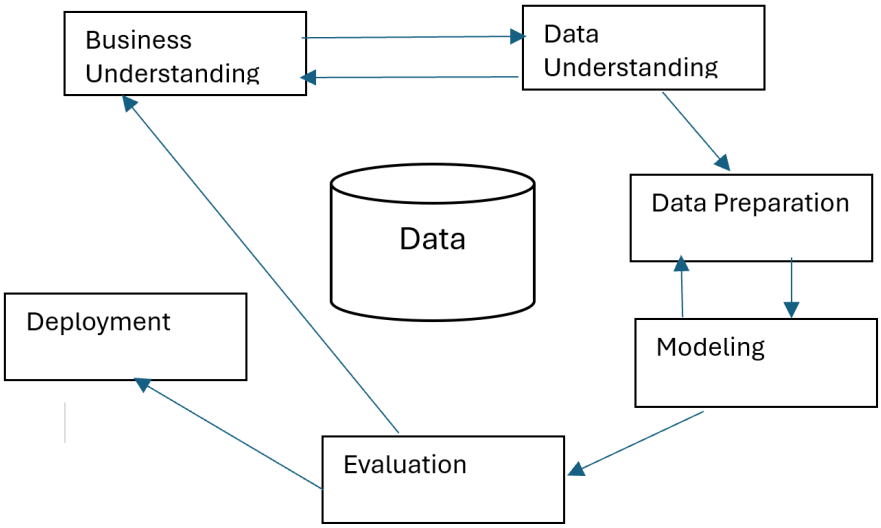

Figure 2. Cross-Industry Standard Process for Data Mining (CRISP-DM)
Recreated by Norman Mooradian

Table 2. Mapping of Paradata Types and Facets

Stage	Paradata Examples	Interpretation	Generation	Purpose
Business Understanding	Project Documents (Project charter, stakeholder registry, project plan	Implicit	Human Made	Incidental
	Ethics Impact Assessments (Algorithmic Impact Assessment (AIA), Data Privacy Impact Assessment (DPIA)	Explicit	Human Made	Incidental, Intentional
Data Understanding	Data provenance record	Explicit	Human Made	Incidental, Intentional
Data Preparation	Documentation (data quality assessment, record of cleansing, corrective steps	Explicit	Human Made	Incidental, Intentional
Modeling	Documentation (algorithm/model type selection)	Explicit	Human Made	Incidental, Intentional
	Documentation (training process, history)	Explicit	Human Made	Incidental, Intentional
Evaluation	Testing outcomes Model scoring	Implicit	Human Made, Automated	Incidental, Intentional
Deployment	Stakeholder communications	Implicit	Human Made	Incidental, Intentional
	User testing	Implicit	Human Made	Incidental
	Incident reporting	Implicit	Human Made	Incidental
	Audit trails	Implicit, Explicit	Automated	Incidental, Intentional
	Audits	Implicit, Explicit	Human Made	Incidental, Intentional

Part 4 - Examples of Paradata elements based on Industry-agnostic and Industry-related use cases (scenarios)

The famous quote, “The future is here—it’s just not evenly distributed,” is attributed to the science fiction author William Gibson. This sentiment applies equally to the use of Artificial Intelligence in organizations for several reasons. We hear about the benefits and drawbacks of AI daily, but the successful implementation of AI is not possible for all organizations for various reasons, including: 1) the lack of resources needed to implement it, 2) the lack of AI literacy necessary to understand its impact within the organization, 3) the fact that AI is a broad umbrella referring to computer systems designed to simulate human intelligence to perform complex tasks including learning, reasoning, problem-solving, perception, and language understanding, and 4) the fact that AI encompasses a wide range of technologies and methodologies to enable machines to perform tasks faster and more reliably than humans.
As stated previously, paradata (documentation of the AI process) will likely be narrowly defined by the process information required for a given context and the AI technologies employed. Although no two use cases will be identical, given the environments in which they operate, the most useful way to understand the concept of paradata in AI processual documentation is to examine AI use cases. As with any process, it is best to take a whole-lifecycle approach to examining the situation and ultimately implementing a solution. The lifecycle illustrated in Figure 3 highlights the high-level stages of an AI process.

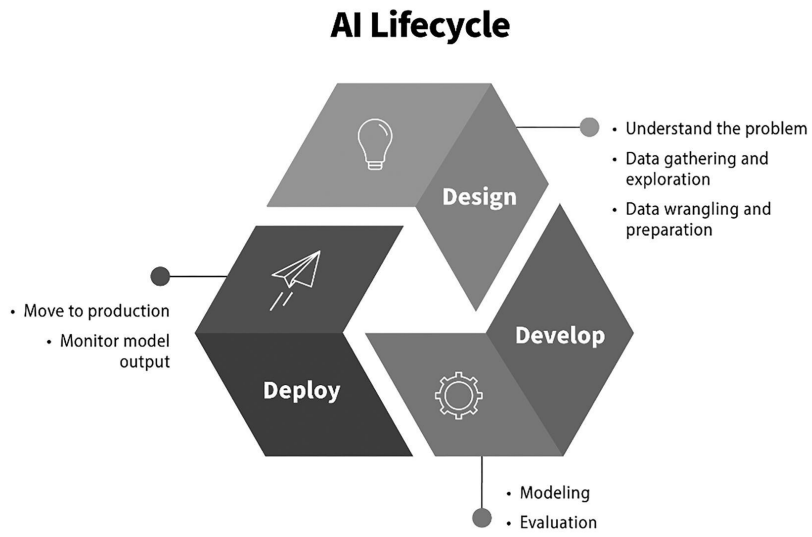

Figure 3. The AI Lifecycle. Source: AI Guide for Government, U.S. General Services Administration, https://coe.gsa.gov/coe/ai-guide-for-government/understanding-managing-ai-lifecycle/

An AI use case begins before an AI solution is selected or developed. In this simplified version (Figure 3), there are three stages: Design, Develop, and Deploy. However, each stage comprises numerous activities and tasks. The first, and likely most important, is to determine if Artificial Intelligence is right for the situation under review. First, the organization must identify the problem or situation they believe AI may resolve, the business challenged faced currently by that problem or situation, potential solutions that may be employed, the benefits to be achieved with each, and then the implementation considerations (e.g., technical, financial, and human), and risk considerations (e.g., confidentiality, security, bias) that must be examined. This phase of the process requires documentation that reveals the factors that informed the decision of whether AI is appropriate for the organization.

Use Case #1: Automated Data Recognition and Extraction for Indexing Digitized Images, Saint Louis Zoo.

At times, it may be determined that AI is not the best solution; however, documentation of the process used to reach this decision will be valuable to those who may question it and propose a similar AI solution in the future. A member of the Paradata team, Rae Lynn Haliday, MBA, CRM, CIGO, FAI, is the Curator of Animal Management Services—Animal Records, Transport and Compliance at the Saint Louis Zoo. The organization currently uses DocuWare for managing animal records.

According to DocuWare's website (https://start.docuware.com/), its *Intelligent Indexing module* employs machine learning to identify key information in documents and transform it into highly structured, actionable data in the form of index terms. Automatic indexing becomes more reliable over time, incorporating indexing corrections made by human operators into its process. In addition to intelligent indexing, DocuWare's *Intelligent Document Processing* capabilities can be purchased at an additional cost. Intelligent Document Processing can classify documents and trigger specific workflows; extract data using OCR; assign extracted data to designated data fields for further processing; and extract handwritten data using Handwritten Text Recognition (HTR).

A summary of the St. Louis Zoo use case is shown in Table 3. Additional details can be found in Use Case #1 in Appendix D.

Table 3. AI Use Case #1: Automated Data Recognition and Extraction for Indexing Digitized Images, St. Louis Zoo.

Category	**Use Case for Automated Data Recognition & Extraction for Indexing Digitized Images of Animal Records from Legacy Microfilm**
Researcher	Rae Lynn Haliday, MBA, CRM, CIGO, FAI Curator, Animal Management Services – Animal Records, Transport & Compliance

Description	The Saint Louis Zoo's mission is to conserve animals and their habitats through animal management, research, recreation, and educational programs that encourage public support and enrich the public experience. This use case is relevant to 40 years of legacy records (created from 1969 to 2009) that were microfilmed and have not yet been digitized. These records constitute four distinct record series and are classified as permanent retention under the zoo-specific records retention schedule. The original camera masters are stored in an off-site microfilm vault; as a local county government, the Zoo qualifies for specific resources, including microfilm storage and digitization services.
Business Challenge	The Saint Louis Zoo's Animal Management Services (AMS) department is working with the State of Missouri Local Records program to digitize the legacy microfilm. The goal of the project is to digitize these records so that this information, with enduring reference, research, and legal value, can be accessed online internally by authorized staff, and the microfilm reader printers and duplicate microfiche can be dispositioned. The current process is time-consuming and training-intensive for interns assisting with the project.
Solution	The use of AI features internal to or integrated with the current EDMS, DocuWare.
Perceived Benefits	AI has the potential to accelerate workflows, including file splitting (extraction) and learning document types for indexing. Once implemented, if effective, it will save employees time, ensure a more consistent approach to digitization, and streamline procedures for interns to assist AMS staff in completing the indexing of legacy animal records.
Implementation Considerations	1. Ability of core AI features within the product to complete the task. 2. The challenges of training the model to work with the organization's files. 3. The cost of additional AI modules, if necessary.
Risk Considerations	1. The organization's content must be protected from unauthorized disclosure. 2. The model must be trained on relevant data to ensure records are retrievable.
Pilot Test:	Testing was conducted in conjunction with the vendor to evaluate features that are part of the current Electronic Document Management System (EDMS) in use for workflows, active, inactive, and legacy animal records, to determine if Artificial Intelligence (AI) applications could be used to speed the work processes, including file splitting (extracting) and learning document types for indexing.
Outcome	Among the challenges identified was that the Intelligent Indexing feature had difficulties reading or interpreting handwritten information due to image quality and poor handwriting. A second option, Intelligent Document Processing, is available from the vendor at an additional cost, but it was impractical due to budget constraints and considering the remaining volume of records that must be indexed at this point in the project. See Use Case #1 in Appendix D for details on the testing process and outcomes.
Final Decision	Ultimately, it was determined that an AI solution was not the best option for this use case. The current process used for digitization, inventory, indexing, and quality control of records identified for this use case is the best path forward, given the human component required for handwritten information, experience with the unique record series, and budgetary constraints.
Reference	Rae Lynn Haliday, Saint Louis Zoo, "Use Case for Automated Data Recognition & Extraction for Indexing Digitized Images of Animal Records from Legacy Microfilm (AI Use Case #1, Appendix D of this report).

Although the decision was made to continue with the current manual process, valuable lessons were learned. Users learned more about the capabilities of their current system and the risks associated with the use of AI when using locally-hosted or cloud-based solutions. As a result,

the organization now has a better understanding of how AI can be used in records management, its risks, benefits, and limitations, and its ability to work with different media (e.g., paper, microfilm, and paper microfilmed and then digitized).

Use Case #2: Bank of Canada, Architecting Accountability in AI

The Bank of Canada started its AI journey by exploring the practical application of Large Language Models (LLMs) for sentiment analysis. Specifically, the Bank used LLMs to assist in the analysis of shifting prices in international oil market data. The Bank combined Bank-developed algorithms and new AI applications to show that the technology could strengthen the Bank's ability to understand changes in commodities markets and apply these capabilities to policy recommendations. As part of this work, we developed a framework for the application of paradata—a key component in ensuring the "explainability" and ethical application of the AI used in formulating market analysis.

Explainability, as found in the GDPR as "The right to an explanation," forms a direct link to understanding AI, the algorithms in use, and their outcomes. For the Bank, explainable AI meant making the tools understandable and interpretable, and this led directly to the deployment of paradata. Paradata documents relationships between AI implementations and their results, making those results understandable in terms of the systems and individuals responsible and accountable for them. One of the Bank's key questions, then, became accountable to whom? Who is the intended audience? In the Bank's case, it's more than one. Ultimately, the Bank is accountable to Canadians, but in terms of immediate reporting, there was a need to provide an explanation to the Managing Director of the Bank's International Economics Department and to the Governor of the Bank. Each of these is a totally different audience and ultimately, accountability rests on principles of authenticity, trustworthiness, and reliability.

The Bank's research for this use case (Table 4) focused on proposing a standard set of paradata and, equally important, what technology could be used to automate the production and maintenance of the paradata.

Table 4. AI Use Case #2: Paradata framework at the Bank of Canada

Category	**Use Case for developing a paradata framework for AI at the Bank of Canada.**
Researcher	Alex Richmond, Enterprise Information Architect, Bank of Canada Mario Beauchamp, AI Data Engineer, Bank of Canada

Description	The Bank of Canada, to meet its ethical AI objectives, developed a paradata framework to support the explainability of its large language models. The use case explored the application of paradata to LLMs used in the support of oil market commodities. The objective of the research was to recommend a minimal set of mandatory paradata that should be created and maintained as part of the development and deployment of LLMs at the Bank. Further, the Bank explored various technologies to automate the production of this paradata.
Business Challenge	The key challenges were as follows: 1. The Bank cannot simply build and deploy an LLM that operates within a black box. The outcomes of the LLM must be explainable. We needed to determine a mechanism to resolve our AI explainability challenges. 2. Determine the minimal set of paradata to be created for each LLM. Paradata covers a wide scope of functionality in AI applications, and the Bank needed to develop a specific mandatory subset. 3. Capturing paradata can be time-consuming. The Bank wanted to explore methods of automating the capture and maintenance of paradata.
Solution	The Bank developed a draft recommendation for paradata to be captured for each LLM base in the ML lifecycle. The Bank's paradata framework recommended four categories of paradata, as follows: ● Training paradata ● Model architecture paradata ● Model evaluation paradata ● User paradata
Perceived Benefits	The application of a standard set of paradata would provide the Bank with the assurance that key aspects of AI explainability would be supported.
Implementation Considerations	Capturing even a bare minimum set of paradata can be time-consuming. Automation is a crucial success criterion. The Bank explored the following tools ● Git LFS (Large File Storage) ● DVC (Data Version Control) ● MLFlow ● Weights & Biases ● JFrog Artifactory
Risk Considerations	The Bank wanted to make sure we could capture more contextual paradata in addition to the pure technical metadata. For example, in terms of explainability, how did economists determine the initial definitions for the labels used in the sentiment analysis of oil-market news articles?
Pilot Test:	The initial pilot test was focused on a limited number of news articles from the Daily Oil Market Bulletin. As the data scientists and economists worked together to develop the LLM, the Bank explored the types of paradata to capture to ensure the explainability of the eventual AI application.
Outcome	A minimal set of paradata was proposed - see the summary above. As of this date, the Bank has not yet formalized the implementation of the paradata framework and is still determining the exact technical components.
Final Decision	Continue the exploration of methods to automate the capture of paradata.
Reference	Document in Appendix D of the RP04 Final Report.

Use Case #3: Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata

While a Paradata Standard may appear to be an attractive solution to the question of what should be captured, how, and for what purpose, the researchers found it to be only part of the answer. Some categories are recommended for all AI use cases, including the use case title, the training dataset, the AI model employed, and the AI Impact Statement. However, other criteria must be considered, including the industry in which it is employed and the purpose for which it is used.

Nancy Powell, SJSU graduate research assistant employed in the Healthcare field; Patricia C. Franks, SJSU Professor Emerita; and Souvick Ghosh, SJSU Assistant Professor, embarked upon a review of seven widely used clinical AI reporting frameworks (CONSORT-AI, SPIRIT-AI, TRIPOD-AI, STARD-AI, DECIDE-AI, MI-CLAIM, MINIMAR) to extract 296 reporting elements and map them to the National Institute of Standards and Technology's Research Data Framework (RDaF Version 2.0) and iterative framework integration. The resulting architecture comprises 113 distinct documentation elements organized across five lifecycle domains and structured as a three-tier architecture spanning governance requirements (59 items), scientific and technical reporting (80 items), and contextual and operational paradata documentation.

Cross-framework analysis revealed a systematic imbalance: existing standards emphasize pre-deployment methodological documentation while underspecifying governance artifacts and post-deployment trace documentation, including workflow integration, human oversight, traceability, and ongoing monitoring. By making these documentation dependencies explicit, this work reframes clinical AI oversight as a problem of information organization and infrastructure, central to auditable, accountable lifecycle governance. An overview of the project is provided in Table 5.

Table 5. AI Lifecycle Documentation in the Healthcare Field.

Category	**Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata**
Researcher	Nancy Powell, GRA, SJSU; Patricia C. Franks, Professor Emerita, SJSU; Souvick Ghosh, Assistant Professor, SJSU
Description	This study asks how documentation–spanning metadata, provenance, and paradata–can be systematically organized across the clinical AI lifecycle to support transparency, governance, and post-deployment oversight. Guided by the National Institute of Standards and Technology's (NIST) Research Data Framework (RDaF) Version 2.0 (Hanisch et al., 2024), we developed a tiered documentation architecture that structures information requirements across regulatory, scientific, and operational dimensions. This architecture is designed to specify the documentation foundation governance requires across the clinical AI lifecycle; it is derived from study-reporting frameworks but reframes their elements as lifecycle documentation artifacts rather than publication requirements.

Business Challenge	Artificial intelligence (AI) systems are increasingly embedded in routine clinical practice, yet lifecycle governance remains difficult to operationalize because the documentation required to support accountability, oversight, and post-deployment monitoring is fragmented and unevenly specified, particularly beyond pre-deployment evaluation.
Solution	An integration and reorganization of reporting elements from existing clinical AI frameworks into a tiered architecture spanning metadata, provenance, and paradata in order to make explicit the documentation dependencies that lifecycle governance implicitly assumes but rarely specifies.
Perceived Benefits	The contribution is an analytic documentation architecture that clarifies where governance-relevant information is expected, absent, or weakly specified across the clinical AI lifecycle.
Implementation Considerations (limitations of the study)	This study has several limitations that reflect its deliberate analytic scope. 5. The analysis focuses on the organization and distribution of documentation requirements encoded in existing reporting frameworks, rather than on how these documentation elements are instantiated, maintained, or used in real-world clinical AI deployments. As such, the proposed tiered architecture is a conceptual scaffolding for documentation rather than an empirically validated implementation model. 6. The framework is derived through integration of established reporting standards that are themselves uneven in scope, maturity, and intended use. Although inter-coder agreement indicates acceptable stability of the integrated Tier 3 schema, variation across frameworks, particularly for contextual and paradata-related elements, highlights the absence of a shared documentation vocabulary for governance-relevant clinical AI information. 7. The comprehensive nature of the Tier 3 base inventory may introduce practical burdens if interpreted as a uniform reporting requirement. The architecture is intentionally non-prescriptive; Tier 3 functions to surface documentation dependencies and gaps, not to imply that all elements must be instantiated in every study or deployment context. 8. The framework has been developed using English-language reporting standards and assumes the availability of documentation infrastructures typical of resource-rich healthcare environments. This may limit generalizability to other contexts and underscores the need for adaptation to diverse regulatory settings, healthcare systems, and sociotechnical conditions.
Risk Considerations	The primary risk concerns the operational burden and “documentation fatigue” that may result from rigidly applying a mandatory checklist rather than a flexible scaffold. The quality of paradata could be sacrificed for the sake of completeness. Semantic ambiguity also represents a risk for inconsistent interpretation.
Pilot Test / Initial Efforts	Initial efforts involved a systematic cross-framework synthesis and metadata mapping of seven prominent clinical AI reporting standards, followed by a three-step refinement process—including NIST RDaF lifecycle alignment and normalization—that distilled 296 raw elements into a consolidated architecture of 113 reporting elements.
Outcome	The clinical AI documentation architecture, a validated, tiered information model bridging the gap between scientific reporting and operational governance, was developed using a standardized lifecycle vocabulary aligned with NIST RDaF 2.0, providing a common language for clinicians, data scientists, and regulatory auditors to manage AI as a continuous information infrastructure.
Final Decision (or Moving Forward)	Future work should empirically examine how such documentation architectures can be instantiated within institutional settings to support auditable, accountable, and adaptive oversight of clinical AI systems over time.

Reference	The researchers have submitted a journal article, “Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata,” for publication.

If we think of paradata as “everything else,” we can understand that some of the 113 documentation items may already be collected as part of the organizational process or the AI system itself; others may not. Flexibility in the use of the term "paradata" is necessary. Figure 4 is an image of the 113 items identified in the study, indicating their inclusion in the three stages of the AI framework.

Figure 4. A portion of the Supplementary Table from the submitted journal article illustrates the lifecycle phase, reporting element/artifact, and tier at which it is specified.

Supplementary Table 1

Architecture of AI Metadata, Provenance, and Paradata Elements Across the Clinical Lifecycle (Clinical AI Data Architecture - CADA - Tier 3)

ID	Lifecycle phase	Reporting element/artifact	Tier 1	Tier 2	Tier 3
4.2	Design Decisions	**AI Model Design**			
4.2.1	Design Decisions	Model description (algorithm type, version, rationale, and development steps)	X	X	X
4.2.2	Design Decisions	Model task (classification, prediction)	X	X	X
4.2.3	Design Decisions	Hyperparameter tuning approach and internal validation (e.g., cross-validation, bootstrapping)	X	X	X
4.2.4	Design Decisions	Comparator models (existing models, human experts)	X	X	X
4.2.5	Design Decisions	Dataset definitions (development/evaluation and training/testing datasets, data origins and formats, baseline characteristics)	X	X	X
4.2.6	Design Decisions	Cohort selection and representativeness rationale		X	X
4.2.7	Design Decisions	Provide prediction formulas, code or implementation details for reproducibility	X	X	X
4.2.8	Design Decisions	Software lifecycle compliance, version control and risk management references	X	X	X
4.2.9	Design Decisions	Paradata: logs and metadata for population, randomization, outcomes, and model development; including traceability controls and cybersecurity safeguards		X	X
5.1	**Data Collection & Processing**	**Operational Oversight & Data Management**			
5.1.1	Data Collection & Processing	Data management, protocol adherence monitoring, and quality assurance		X	X
5.1.2	Data Collection & Processing	Data sources described (e.g., sensors, surveys, laboratory tests, transaction records, digital platforms)	X	X	X
5.1.3	Data Collection & Processing	Reliability and validity of instruments/tools used for data collections		X	X
5.1.4	Data Collection & Processing	Data collection forms, schemas and protocols			X
5.1.5	Data Collection & Processing	Recruitment/selection and follow-up details (where/when identified, timeframes, retention strategies)			X
5.1.6	Data Collection & Processing	Discontinuation/deviation handling		X	X
5.1.7	Data Collection & Processing	High-level participant flow and follow-up			X
5.2	Data Collection & Processing	**AI Input Data Preparation**			
5.2.1	Data Collection & Processing	Acquisition and selection of input data	X		X
5.2.2	Data Collection & Processing	Pre-processing and transformations of data applied (e.g., normalization, feature extraction, augmentation)	X	X	X
5.2.3	Data Collection & Processing	Planned data quality checks (missing data, poor quality data, reasons for omitting data)	X	X	X
5.2.4	Data Collection & Processing	Missing data handling	X	X	X
5.2.5	Data Collection & Processing	Security and storage procedures	X	X	X

Any of the items may be considered paradata, dependent upon the organization’s approach to AI documentation. However, note line 4.2.9, which includes specific elements of paradata not included elsewhere in the table. While this study identifies potential paradata elements for collection; the specific elements collected will vary depending on the AI application and purpose, as well as governing laws and regulations.

Use Case #4: CUO5: Information Framework for Documenting Case Studies

CU05 is a companion study under the umbrella of InterPARES Trust AI led by Maria Guercio, and Stefano Allegrezza. The official title is "The role of AI in identifying or reconstituting archival aggregations of digital records and enriching metadata schemas." The study was based on the premise that "the uncontrolled creation of a huge numbers of current records with missing metadata, necessary to ensure the reliability, trustworthiness, quality, and sustainability of appraisal and acquisition, is a common and complex problem today" (https://interparestrustai.org/trust/about_research/studies).

As part of the CU05 research project, literature on paradata and additional information elements was analyzed to determine the information required for documenting the CU05 case studies developed by that team. Rather than base their analysis on a commercially available tool such as Google Model Cards or IBM AI Fact Sheets, the team analyzed the steps they followed in the AI process post-hoc. A brief summary of the Use Case is presented in Table 6. Refer to the Final Project Report for additional information on the InterPARES CU05 Project: Information Framework for Documenting Case Studies, along with an image of the excel spreadsheet illustrating the seven stages of the AI Lifecycle defined and three phases of data collection: Planning, Executing and Closing.

Table 6. AI Use Case #4 – InterPARES CU05 Project: Information Framework for Documenting Case Studies

Industry: Archive and Records Management

Category	**InterPARES CU05 Project: Information Framework for Documenting Case Studies**
Researchers	Maria Guercio, Stefano Allegrezza, Francesca Magnoni, Grandi Massimiliano, Bruna La Sorda, Maria Mata Caravaca
Description	In the CU05 research project, literature on paradata and additional information elements was analyzed to determine the information required for documenting the CU05 case studies. The literature review led to the establishment of a framework that combines two primary resources: the Machine Learning life cycle (Scott Cameron, 2024), which encompasses the process activities involved in developing and deploying AI models and the phases of data collection based on the CCSDS recommendations IPELTU and its category of additional information (IPELTU 2024), which outlines (in compliance with the OAIS standard) the stages of complex projects that generate critical information throughout the activities involved in implementing AI.
Business Challenge	The advent of AI presents the need for the collection of additional information elements that provide insights into the process and methodology used, as well as the decisions and challenges faced throughout the case studies. This documentation is crucial for evaluating the effectiveness and impact of AI integration into archival practices.

Solution	Although planning the creation and collection of paradata in the beginning to automate as much as possible is ideal, this study took an ex-post approach to document the actions taken and decisions made related to the machine learning process followed. To do so, the team identified, collected, and listed the up-to-date available material from the case study, which includes the initial project proposal, a log of the records used for the study, the logs produced by the machine learning model exercises, when available, and the machine learning report written by RecordPoint (the system vendor). The aim is to map the information we have collected from the case study against the identified AI project phases and stages. In addition, we are drafting the various stages of the project with the available information and its processes in a descriptive form.
Perceived Benefits	This study provides an example of a paradata framework designed for a specific use case (analyzing case studies). The contribution is one that can be developed further by considering how some of the system paradata creation and collection can be collected automatically. It provides an example of how the archive can work with a vendor to understand AI processual documentation. And it can provide the basis for other departments and institutions to begin to understand where governance-relevant information is expected, absent, or weakly specified across the clinical AI lifecycle.
Implementation Considerations (limitations of the study)	This study has several limitations that reflect its deliberate analytic scope. 1. The analysis focuses on the organization and distribution of documentation requirements developed for one archival Use Case. 2. The framework is based upon paradata information gathered through articles up to 2024-—this field is developing rapidly. 3. The framework has been developed using CCSDS recommendations IPELTU and its category of additional information (IPELTU 2024), which outlines (in compliance with the OAIS standard). Based on the industry in which the organization operates, other standards and regulatory requirements may be more appropriate.
Risk Considerations	The approach taken is ex-post, which means some actions and decisions may have been missed when documenting the AI process. The primary risk involves “documentation fatigue,” which may be alleviated through the integration of the automated creation and collection of paradata within the system itself.
Pilot Test / Initial Efforts	Initial efforts produced a documentation framework for the collection of paradata based on the AI lifecycle and existing archival standards. In parallel to the above-described framework, the team applied schema to the *NATO Archives-RecordPoint case study*. The aim was to see how an AI project can be documented from the perspective of the practitioners, and to test the framework in a real case scenario. The perspective is the one of an archivist that is trying to document these processes. A description of the NATO Archives-RecordPoint case study will be provided in the CU05 Final Report (February 2026).
Outcome	The documentation is detailed across the AI application lifecycle and the phases of data collection from: IPELTU 2024. The framework can provide a useful starting point for other departments and organizations seeking a way to document the actions and decisions made when implementing an AI tool within the archival field.
Final Decision (or Moving Forward)	Future work should empirically examine how this framework can be employed to support auditable, accountable, and adaptive oversight of AI systems over time.
Reference	See CU05 Final Report, February 2026.

The amount of time and effort that went into this study indicates the new skills and knowledge demanded of today’s archival profession. Part 5 includes reflections of research members on

the concept of paradata, and the impact AI will have on the archives profession and professionals.

Part 5 - Reflections of PARADATA Team Members

Members of the PARADATA Research team were asked to share their thoughts on paradata, both conceptually and in practice, as a result of their involvement in the study, *Preserving AI Techniques as Paradata*. The complete responses provided to four prompts are included in *Appendix C - Reflections in the Words of Researchers*. This section of the report provides a summary and analysis of those reflections.

5.1 What can we conclude as a result of the work of the RP04 study?

A consolidated set of **10 topical headings** synthesizes the points provided by researchers who responded to this question, while preserving their conceptual distinctions and emphases:

1. **Paradata is a Distinct Conceptual Framework**

   Paradata is conceptually distinct from metadata, provenance, and explainability (XAI). While related, paradata refers specifically to the documentation of processes, decisions, transformations, and contextual factors shaping AI systems. It serves as both:

   - A category of information, and
   - A heuristic tool for identifying documentation needs across the AI Lifecycle.

   This conceptual clarification of paradata is one of our study's core contributions to ITrustAI.

4. **Paradata is Essential to AI Accountability, Trust, and Ethical AI Governance**

   Paradata (alongside metadata) contributes to accountability, fairness, trust, and responsible AI Governance. Paradata is fundamental to:

   - Transparency
   - Accountability
   - Fairness
   - Understandability
   - Trust in AI systems

Because AI processes are often opaque ("Blackbox"), documenting lifecycle decisions and system behaviors become essential for governance and public confidence.

3. **Complementarity of Metadata and Paradata**

   This study suggests that metadata and paradata are complementary components of responsible AI documentation. Each plays a distinct role in describing datasets, processes, decisions, and context.

   - Metadata structures and contextualizes datasets.
   - Paradata captures decision-making processes and operational actions.

   Together, they support ethical AI governance, yet paradata remains less standardized than metadata frameworks.

4. **The Challenge of Defining the "AI Record"**

   A major archival implication is the need to define and construct an **AI record**:

   - A record of decisions and/or actions taken by AI systems.
   - Evidence tied to rights, interests, and impacts on individuals and groups.

   There is much work yet to be done in defining, identifying, and constructing "AI records." Because AI decision-making is frequently opaque, documenting the reasoning is difficult. Paradata provides a conceptual mechanism for identifying the information necessary to build such records.

5. **Legal, Regulatory, and Governance Drivers**

   The study findings highlight how laws, standards, and governance frameworks:

   - Explicitly require recordkeeping and documentation, or
   - Implicitly imply documentation requirements.

   Examples referenced during our study include the EU AI Act, the NIST AI Risk Management Framework, and domain-specific standards such as STARD-AI. However, these frameworks often lack specificity. Paradata helps interpret and operationalize these expectations for documentation and recordkeeping practices.

6. Paradata Across the AI Lifecycle

Identification of information, decisions, processes, and activities generated throughout the AI lifecycle is essential for records creation. This study emphasizes that documentation must span the entire AI lifecycle, including:

- Data collection
- Data transformation
- Model development
- Deployment
- Monitoring
- Digitization processes

Retrofitting documentation is inefficient; instead, paradata must be embedded in workflows from inception and linked to lifecycle artifacts in a structured manner.

7. Paradata in AI-Assisted Digitization and Legacy Records

The study extends beyond AI system development to AI-assisted digitization of legacy records, highlighting:

- Quality control challenges (e.g., OCR accuracy, image quality issues)
- Risks in indexing handwritten or degraded records
- The need for documenting AI performance limitations

Paradata can help identify errors, biases, and quality deficiencies while supporting long-term preservation and integrity of digitized records.

8. Paradata as a mechanism for Juridical and Archival Compliance

Paradata refers to data about the processes, methodologies, and technical actions involved in creating, managing, or analyzing data, and serves as crucial documentation to ensure both archival accountability and juridical compliance. While traditional metadata describes the content of a record, paradata describes the *context* and *decisions* behind it, which is essential for proving the authenticity, integrity, and provenance of records in legal, archival, and AI contexts. Examples of paradata for juridical/legal and archival compliance include:

- ***Paradata in Archival Compliance—Procedural Integrity***: Paradata records the systematic organization and, for instance, in digitization projects, the specific steps taken to ensure the digital reproduction is authentic.
- ***Paradata in Juridical/Legal Compliance—Legal Hold and E-Discovery***: Paradata assists in locating specific records and demonstrating that a "tamper-proof" retention policy has been followed.

9. Operational and Cultural Implementation Challenges

Practical lessons from institutional implementation include:

- Manual capture of paradata is costly; automation is essential.
- Governance teams and developers use different terminologies.
- Shared definitions and taxonomies are critical.
- Paradata must be structured and linked to artifacts to be useful.

Successful implementation requires not only technical design but also cultural adoption and organizational alignment.

10. Extending Paradata Beyond AI

A key conclusion is that the need for paradata is not limited to systems labeled "AI." Many technical and cultural technologies raise similar documentation challenges. The relevance of paradata offers opportunities for collaboration among records professionals, technologists, and other stakeholders. Thus, this study suggests that paradata may represent a broader records management strategy for documenting complex socio-technical systems in general.

Overall Conclusions of Researchers' Responses to Question 1: What can we conclude as a result of the work of the RP04 study?

The RP04 study concludes that paradata is both conceptually necessary and practically indispensable for documenting AI systems and related technologies. It provides a framework for identifying, structuring, and preserving information required to meet accountability, legal, archival, and governance expectations. Moreover, it repositions documentation as an embedded lifecycle function rather than an after-the-fact compliance exercise.

5.2 What one change/action does the work produced as a result of the RP04 study suggest individual archival and records management practitioners should be making now?

A consolidated set of **10 topical headings** integrates and synthesizes the perspectives of researchers who responded to this question, while keeping the focus on actionable implications for archives and records management:

1. **Prioritize the Capture and Preservation of Paradata**
   Practitioners must actively ensure that paradata about AI systems and AI-enabled processes is created, captured, and retained. Paradata is essential for documenting how records are generated, transformed, or influenced by AI systems and must be treated as integral to recordkeeping rather than incidental.

2. **Evolving Concepts of Records in AI Contexts**

   The notion of a "record" is both persisting and evolving. Records and information professionals need to understand how paradata fits within evolving understandings of what constitutes a record. AI systems produce or influence records in ways that require expanding traditional definitions to include metadata and paradata, particularly where AI opacity challenges evidentiary value.

3. **AI Literacy and Professional Education**

   AI literacy is now a core competency. Archivists and records managers must understand:

   - How AI tools function at a system level
   - How decisions are made within AI processes
   - What documentation should exist at each stage of the AI lifecycle

   Building this literacy enables practitioners to critically assess AI systems rather than passively adopt them.

4. **Interrogate AI Tools Before, During, and After Use**

   Practitioners should routinely ask:

   - What paradata does this tool generate?
   - What explanations are available?
   - What decision trails exist?

   This interrogation should occur at procurement, implementation, operation, and review stages of the AI lifecycle.

5. **Make Paradata a Core Procurement Requirement**

   For off-the-shelf and SaaS AI systems, the availability, quality, and usefulness of paradata should be a primary evaluation factor. Archivists must be able to compare vendors not only on functionality but also on transparency, explainability, and documentation support.

6. **Align Paradata Requirements with Legal and Juridical Obligations**

   Paradata requirements must be understood in light of legal mandates, regulatory frameworks, and preservation needs. Paradata design must be driven by:

   - Legal mandates
   - Regulatory requirements
   - Chain of Preservation (COP) obligations
   - Institutional functions

   Documentation requirements should be derived from juridical and governance frameworks and translated into specific archival processes.

7. **Recognize and Incorporate Human-Created and Machine-Generated Paradata**

   Paradata exists in multiple forms:

   - *Human-generated* (policies, communications, training documentation)
   - *Machine-generated* (logs, decision traces, model outputs, system metadata)

   Each type serves different evidentiary purposes, and both must be incorporated into AI records.

8. **Identify Implied Documentation in Compliance and Governance Mandates**

   Even when regulations do not explicitly define record requirements, activities such as testing, training, validation, documentation, and communication imply the existence of paradata. Practitioners should proactively identify and capture these data trails as components of AI records.

9. **Address Transparency, Accountability, and Risk in AI Systems**

   Paradata is a tool for capturing both system-level explanations and records of specific AI decisions. It is essential to:

   - Mitigate opacity in AI systems
   - Counter risks such as model hallucination
   - Support accountability and trust
   - Define levels of certainty and threat models

   Documentation should balance transparency with security, particularly for sensitive systems.

10. **Operationalize and Sustain Paradata Practices**

    Paradata is a unifying framework to bring together diverse forms of information required to support AI governance, accountability, and archival practice. To make paradata viable long-term, institutions must:

    - Harmonize definitions across disciplines
    - Clarify responsibility for recording and verifying paradata

- Ensure quality control
- Capture paradata in real-time or high-velocity environments
- Embed paradata requirements into governance frameworks

Sustainable AI documentation requires institutionalization, not ad hoc efforts.

Overall Conclusions of Researchers' Responses to Question 2: What one change/action does the work produced as a result of the RP04 study suggest individual archival and records management practitioners should be making now?

In sum, this study suggests a clear change in practice: archivists and records managers must treat AI documentation and paradata capture as foundational recordkeeping work. This requires building AI literacy, embedding paradata requirements into procurement and governance, expanding the concept of records, and establishing sustainable documentation responsibilities across the AI lifecycle.

5.3 Has/How has the work produced as a result of the RP04 study changed your thinking on how documentation of the AI process can best be accomplished?

A consolidated set of **10 topical headings** synthesizes the perspectives of researchers who responded to this question and captures both the conceptual and practical dimensions of paradata in AI contexts:

1. **Limits of Generic Paradata Standards**

   Generic standards alone are insufficient and must be grounded in domain-specific AI uses and contexts. Keep in mind:

   - The initial expectation of a “nice and tidy” paradata standard has shifted.
   - Standards are necessary but insufficient on their own.
   - There is a risk of prematurely developing generic AI paradata that lacks contextual grounding.
   - Emphasis must be placed on avoiding overly abstract or decontextualized frameworks.

   Core Shift: There is a shift from reliance on a universal paradata standard to recognition of the complexity and contextual specificity of AI documentation.

2. **Contextual and Domain-Specific Understanding**

   Context-sensitive implementations are necessary for the following reasons:

- Documentation must be grounded in specific domains and systems.
- Archival Concepts such as accountability must be reinterpreted within AI contexts.
- AI use varies significantly across settings; documentation must reflect this.

Core Shift: There is a shift from a generic conceptualization of the role of paradata in documentation to context-sensitive implementation, which also implies industry-based contexts.

3. AI as a Sociotechnical Process (not just a tool)

AI processes are sociotechnical systems that share similarities with non-AI decision-making and record-keeping processes. Some points to consider:

- *Earlier framing*: documenting a technology/tool.
- *Revised view*: AI processes resemble broader sociotechnical processes.
- *Focus on training data, evaluation, and human action r*ather than AI as an autonomous entity.
- An effort must be made to *resist anthropomorphizing AI*.

Core Shift: A shift from technological determinism to sociotechnical integration is occurring.

4. Paradata as part of an Archival Ecosystem

There is a need to reinterpret foundational archival concepts—such as accountability and evidence—within AI-powered processes.

- Paradata complements and is complemented by other records.
- AI documentation exists alongside administrative and organizational records.
- Paradata itself may be an “archive.”

Core Shift: A shift from isolated documentation artifacts to an interconnected archival network is called for.

9. Automation of Paradata Capture

Automating paradata creation within AI systems is necessary to enable traceability and auditability across the AI lifecycle.

- Manual documentation is impractical.
- Paradata must be embedded in AI systems.
- Real-time capture ensures traceability and auditability across the AI lifecycle.

- Paradata must be integrated into AI workflows where possible.

Core Shift: A shift from manual, retrospective documentation to automated, system-integrated capture is necessary.

10. Integration of Metadata and Paradata into AI Workflows & Documentation Practices

Both metadata and paradata must be more effectively integrated into AI workflows and documentation practices. For example,

- There is a need for a better understanding of how metadata and paradata interact.
- Documentation must be built into the lifecycle rather than appended afterward.
- A lifecycle approach must be taken to facilitate traceability.

Core Shift: There is a shift from external documentation practices to embedded lifecycle governance

11. Aggregating Distributed Paradata and Documentation

Both challenges and solutions exist for aggregating paradata and related documentation across systems, formats, and organizational units.

- Supporting documentation exists in multiple forms (policies, training records, impact statements, vendor documentation).
- There is a need for centralized aggregation.
- A portal/inventory of AI projects and required documentation should be developed.
- This documentation must serve multiple stakeholders, but not all will require the same materials depending on their role in the AI process.

Core Shift: A transition from dispersed documentation to a centralized governance architecture is occurring.

12. Human and Organizational Dimensions of Paradata

Human decision-making, governance, and organizational practices are critical components of paradata and warrant emphasis.

- There is a shift away from technical bias.
- Human decisions are often more important than system logs.
- Key documentation questions are organizational:

- Who selected the tool?
- How was it adapted?
- What Quality Assurance measures were taken?
- What monitoring occurred?
- How are stakeholders informed?

Core Shift: A shift from a focus on technical logging to organizational accountability is necessary.

9. Technical Logs are Necessary but Insufficient

Complementary but distinct roles of automated technical logs and human-generated contextual records are central to AI accountability.

- Technical logs remain important artifacts.
- However, logs alone cannot ensure accountability.
- AI Blackbox issues are real, but governance questions are more pressing.

Core Shift: Priorities for technical and organizational documentation must be rebalanced.

10. Paradata as a Mitigation Strategy Under Constraints

Using selectively generated paradata helps the organization manage risk, uncertainty, and resource constraints while enabling responsible use of AI.

- Paradata can be described as “everything else” needed to enable examination.
- Ideally, practice should preserve the software, mechanisms, and data necessary for reproducibility.
- Real-world constraints require selective, strategic generation of paradata.
- AI’s cost/performance advantages must be balanced against uncertainty.

Core Shift: A shift from exhaustive preservation ideals toward pragmatic, constraint-aware documentation strategies is underway.

Overall Conclusions of Researchers' Responses to Question 3: Has/How has the work produced as a result of the RP04 study changed your thinking on how documentation of the AI process can best be accomplished?

The members of the research team have undergone a shared evolution in their thinking regarding paradata, as shown in Table 7.

Table 7. Early Assumptions and Revised Understanding of Paradata for the AI Process

Early Assumptions	**Revised Understanding**
Standardization might solve the problem	Contextual, domain-specific work is essential
AI documentation is mainly technical	Organizational and human factors are central
Manual documentation is feasible	Automation and workflow embedding are required
AI is primarily a tool	AI processes are sociotechnical systems
Paradata is supplementary	Paradata is foundational to accountability and preservation.

5.4 How has the work produced as a result of the RP04 study changed the way you think we should envision the role of the archivist/records manager for the future?

A consolidated set of **10 topical headings** synthesizes and analyzes the points provided by researchers who responded to this question, with attention paid to professional identity, skills, and leadership roles:

1. **From Managing Objects to Governing Systems**

   Archivists must shift from managing static records to overseeing dynamic technical and sociotechnical systems that generate records.

2. **Documentary Accountability as Core Expertise**

   The future archivist becomes a specialist in ensuring accountability, traceability, and evidentiary integrity within automated environments.

3. **Centrality of Paradata**

   Paradata—documentation about how records are created, processed, and transformed—emerges as essential to future recordkeeping practice.

4. **Standardization of AI Paradata**

   AI-related documentation must be formalized into standards where possible, integrated alongside metadata requirements, and embedded in system design. However, Domain-Specific paradata is also necessary—for example, Paradata in Healthcare and Banking.

5. **Technological and Computational Literacy**

   Archivists and records managers must deepen their fluency in computer science concepts, digital infrastructure, and automated systems to advocate effectively for archival principles.

6. **Ethical Stewardship of AI**

   The profession has a responsibility to model grounded, critical, and transparent AI implementation, including the documentation of limitations and the mitigation of bias.

7. **Interdisciplinary Collaboration**

   Archivists must actively engage in interdisciplinary teams—working with technologists, policymakers, researchers, and system designers.

8. **Operationalizing Theory into Practice**

   Conceptual frameworks (e.g., paradata) must be translated into operational procedures, workflows, and governance structures.

9. **Balancing Specialization and Scope**

   The role may "contract" toward deeper technological specialization while simultaneously "expanding" into broader collaborative and translational functions.

10. **Communication and Translation Across Audiences**

    Archivists serve as bridges—articulating archival needs to AI developers, refining technical concepts for users, and explaining the impacts of automation to scholars and society.

Core Shift: The future role for archivists and records managers is clearly shifting toward AI-aware, system-centered, accountability-focused recordkeeping.

Overall Insight

The RP04 study positions archivists and records managers not as passive custodians of digital outputs, but as active designers, regulators, translators, and ethical stewards of AI-enabled recordkeeping ecosystems. The profession's future lies at the intersection of accountability, technology, and interdisciplinary governance.

A list of AI and Paradata topics for inclusion in archives and records management training and education is provided in Appendix E.

Part 6: Reference chart for assessing paradata needs

Table 8, shown below, is also available as a spreadsheet as linked: [paradata reference chart](#).

While the specific context of a given AI application remains the key factor in assessing the necessary paradata, general outlines of common paradata to be retained for transparency may be prepared since the general outlines of the AI process remains similar across contexts. The following table highlights categories of information which comprise relevant paradata across multiple contexts, and details a few typical artefacts which are commonly required for preservation. Rather than providing a one-size-fits-all checklist, the table is intended to provide the basis for an analysis of the relevant paradata for a given application context, which must be built out and detailed based on a contextually-situated analysis. Each context will have specific organizational and ethical risks associated with AI implementation which must be carefully attended to. Part of these risks may be analyzed using the traditional tools of organizational risk analysis, while the ethics of the archival profession can and should provide another. InterPARES has produced guidance to understanding critical archival perspectives in relation to artificial intelligence and machine learning tools (see Fewster, K. and R. Arias-Hernandez, 2024). AI implementers in archives should adapt their solutions to the cultural and social context in which they operate, rather than taking a predetermined, one-size-fits-all approach which is insensitive to particularity.

Table 8. Paradata Elements Reference Table which includes questions paradata may resolve.

Paradata elements reference table		
Category	**Short title**	**Questions paradata may need to resolve**
Information about the application context	Risks of operating context	What are the specific risks of the operating context? What risks are amplified or diminished when introducing an AI tool? How were these risks assessed and is the risk assessment available?
	Legal context	What are the legal requirements of this operating context?
Information about the model	AI tool itself	What AI model or tool was used? Can this tool be preserved?
	Risk analysis: AI tool	How was the model's suitability assessed for this application? How is it suited or adaptable to the risks emergent in this application?
	Setup	What parameters were used in the model's setup, training, and execution? Are there any alterations made from its stock setup? What purposes did these alterations serve?

Paradata elements reference table		
Category	**Short title**	**Questions paradata may need to resolve**
	Limitations to replicability	Does the model introduce randomness into the process? What inherent obstacles to replicability exist?
Information about the training dataset	Training data	What training dataset was used? Where is it accessible or preserved?
	Risk analysis: training data	How was the training dataset's suitability assessed for this application? What limitations or qualifications does this dataset have? What risks may emerge from using this dataset in this context?
Information about quality assurance	Necessary level of certainty	What level of certainty is needed for this application? How will quality be assessed before implementation?
	Certainty provisions of AI tool	How does the model qualify uncertain conclusions? Does it identify sources of uncertainty?
	Ongoing monitoring	How is the tool's use monitored over time? How can implementers monitor its actions to assure its ongoing efficacy and avoid doing harm?
Processual documentation or logs generated by AI tools	Operating transparency	On what basis has the model make particular decisions? What explanations are available for the individual decisions of the AI system?
	Understandability of outputs	How are the records processed by this AI system's actions tracked? Are these relationships clearly recorded for administrators and researchers for collections management and transparency purposes?
Information about rights of data subjects	Right to information	What right to information do data subjects or citizens possess? Can individuals request data or information? Is the model's implementer obligated to inform data subjects? What information from the AI process can/should be made public?
	Right to privacy	What rights to privacy do data subjects or citizens possess?
	Right to redress	What avenues do data subjects and those effected by the AI system have for redress of harms caused by it?

Part 7 - Conclusions and Recommendations

This project was initiated to answer a fundamental question posed by the use of Artificial Intelligence: "If an AI technique is used to facilitate or automate an archival, recordkeeping, or

other process, how much of that AI technique, its code, the data (probably a subset of existing records) we use to train it, test cases and test results to examine its efficacy, its parameters and their values at or over the time of application, the technical environment in which it is executive, and the records it (the AI technique) is applied to for automation purposes, should be preserved?"

The archivists, records managers, information professionals, and AI developers comprising the Paradata team looked to current and emerging regulations for guidance. Although most documents, such as the EU AI Act and the NIST AI Risk Management Framework, indicated implicit or explicit documentation requirements, none addressed the gap that exists between many of those requirements and the traditional forms of documentation: metadata, Explainable AI, and records. The study team recognized the need for a term to describe the "everything else" that must be created, captured, and preserved and agreed that the term Paradata, already used in cultural heritage, archeology, and statistics, could be repurposed for our needs. The definition of the term revised to document the AI process, was approved:

> **Paradata** (n.) Information about the procedure(s) and tools used to create and process information resources, along with information about the persons carrying out those procedures (ITrustAI terminology, n.d.).

What appeared from the outset to be a straightforward task of developing standards for Paradata for the AI Process became a four-year journey that provided multiple lessons learned and left the team with many avenues for future use.

The main conclusions agreed upon by researchers by the end of this study are:

1. Paradata is both conceptually necessary and practically indispensable for comprehensively documenting AI systems and related technologies. It provides a framework for identifying, structuring, and preserving information required to meet accountability, legal, archival, and governance expectations. Moreover, it repositions documentation as an embedded lifecycle function rather than an after-the-fact compliance exercise.
2. A clear change in practice is required of archivists and records managers who must treat AI documentation and paradata capture as foundational recordkeeping work. This requires building AI literacy, embedding paradata requirements into procurement and governance, expanding the concept of records, and establishing sustainable documentation responsibilities throughout the AI lifecycle.
3. The members of the research team have undergone a shared evolution in their thinking regarding the use of paradata to document the AI process to emphasize the following factors: Contextual, domain-specific work is essential; organizational and human factors are central; automation and workflow embedding are required; AI processes are sociotechnical systems; and Paradata is foundational to accountability and preservation.

4. Archivists and records managers are not passive custodians of digital outputs but active designers, regulators, translators, and ethical stewards of AI-enabled recordkeeping ecosystems. The profession's future lies at the intersection of accountability, technology, and interdisciplinary governance.

As illustrated by the literature review completed, guidance documents studied, and use cases developed, there is no one easy answer to what paradata must be created, captured, and preserved in order to document the AI process. AI processes are situational and dependent not only upon the technology employed and the environment in which the organization operates but also the humans tasked with its governance. Further research can be invested in the following activities:

1. Development of a "core" set of Paradata standards for AI processes drawing on current and emerging regulations and AI guidance available.
2. Development of Paradata templates for AI processes in various industries and use contexts—perhaps starting with the Use Case included in this report for the Healthcare field.
3. Collaboration with vendors of enterprise content management systems to determine what paradata can be automatically created, captured, and preserved for as long as necessary.
4. Additional work in the area of automated AI Governance to optimize costs and manage risks related to AI implementation.
5. Development of methodologies to capture and preserve paradata while maintaining the necessary balance between the right to know of the public/regulators and the rights to privacy of individuals and populations concerned.
6. Refinement of the AI and Paradata recommendations for inclusion in archives and records management education and training presented in Appendix E.

References

Cameron, S., P. Franks, and B. Hamidzadeh. 2023. Positioning Paradata: A Conceptual Frame for AI Processual Documentation in Archives and Recordkeeping Contexts. Journal of Computing and Cultural Heritage 16:4, Article 75 (December 2023). https://doi.org/10.1145/3594728.

Cameron S., P. Franks, I. Huvila and N. Mooradian. 2025. Navigating accountability: the role of paradata in AI documentation and governance. *Journal of Documentation* 81 (4): 906–926, doi: https://doi.org/10.1108/JD-01-2025-0009.

Fewster, K. and R. Arias-Hernandez. 2024. Critical AI/ML and AI/ML Ethics for Archival Professionals. InterPARES Trust AI.. https://interparestrustai.org/assets/public/dissemination/AD_01_Module2_v_1_0_20241211.pdf

Shinde, G.; T. Kirstein; S. Ghosh and P. Franks (2025, October), “Tracing the Past, Predicting the Future: A Systematic Review of AI in Archival Science,” Proceedings of the Association for Information Science and Technology 62 (1): 658-671.

Appendix A
Contributing Members

Mario Beauchamp is a Computer Scientist, AI Scientist, and Data Scientist, possessing experience in software architecture, software design, and research centered on Neural Networks models. Following an 11-year career in the public sector, they have transitioned to the private sector to advance the development and research of AI agents, particularly focusing on advancements in deterministic AI and long-context scenarios.

Scott Cameron is an archivist and librarian. He has contributed to the InterPARES research group on paradata since 2022, and presently works in research data management at the Bank of Canada. Scott holds an MA in History from the University of Toronto and library and archival degrees from the University of British Columbia. Responsibility for the views expressed should not be attributed to his employer.

Fred Cohen is an industry analyst, consultant, trusted advisor, digital forensic evidence examiner, educator, and author of hundreds of articles. He is a Fellow of ISC2, CEO of Management Analytics, Managing Member of Angel to Exit, and on the editorial boards of several publications. His current research interests involve cybersecurity, artificial intelligence, and archives and records management.

Jeremy Davet, a former graduate research assistant, received his Master of Archival Studies from the University of British Columbia in 2022. He is currently the Digital Preservation Librarian at the University of Houston Libraries, responsible for the ongoing development of the library's digital preservation strategy and infrastructure. He also serves on the Society of Southwestern Archivists (SSA) Ad Hoc Committee on Sustainability.

Kaila Fewster is an archivist and librarian, holding MAS and MLIS degrees from the University of British Columbia. She worked as an InterPARES researcher to develop educational modules for archival professionals using Al and investigate the role of paradata in capturing Al records. She has also worked in conservation roles at UBC and the Roman Catholic Archdiocese of Vancouver.

Patricia C Franks is a Certified Archivist, Certified Records Manager, and Information Governance Professional. She is a past president of the National Association of Government Archivists and Records Administrators and author, editor, or co-editor of numerous publications, including *Records and Information Management* and *The Handbook of Archival Practice*. Her current research interests involve artificial intelligence and archives and records management.

Rae Lynn Haliday, CRM, CIGO, Curator of Animal Management Services—Animal Records, Transport and Compliance at the St. Louis Zoo, is a past President of the Institute of Certified Records Managers and a frequent speaker and author. She teaches records management, leadership, and advocacy for records professionals and currently serves as a Director on the ARMA International Board of Directors.

Babak Hamidzadeh served as Interim and Associate Dean at the University of Maryland Libraries. He was previously Director of the Repository Development Center at the Library of Congress, held senior roles at Boeing, and taught at the University of British Columbia and Hong Kong University of Science and Technology. He earned his doctorate in computer science from the University of Minnesota.

Kat Hodgson is an archivist, librarian, and linguist with MAS and MLIS degrees from the University of British Columbia. Currently, she works as an Assistant Archivist for the Whistler Museum and Archives Society, where she is cataloguing one of their largest photographic collections. Previously, she has worked as a student cataloguing librarian at Library and Archives Canada.

Isto Huvila is a chaired professor of information studies at the Department of ALM at Uppsala University in Sweden. His primary areas of research include information, knowledge and records management, information work, knowledge organization, documentation, and social and participatory information practices. Huvila has published extensively on paradata, including two books published by Springer and Cambridge University Press.

Norman Mooradian is an Assistant Professor in the School of Information at San Jose State University. He has published articles on information ethics, business ethics, information privacy, enterprise content management (ECM), knowledge management, and virtual reality, and is the author of Ethics for Records and Information Management (2018, ALA). Mooradian teaches a seminar titled AI, Data, and Ethics, as well as courses in information science.

Souvick Ghosh is an Assistant Professor at the School of Information, San José State University. His research focuses on Applied, Ethical, and Conversational AI, with emphasis on improving the lives of marginalized communities through AI applications. Since 2020, Dr. Ghosh has published four peer-reviewed journal articles, 26 peer-reviewed conference papers in proceedings, one book (in progress), and two book chapters.

Nancy Powell is a graduate student in the Master of Science in Informatics program at San José State University, where she contributes to research on AI reporting frameworks for clinical use. She also works as a clinical quality analyst, bringing healthcare experience to her academic and research pursuits.

Alex Richmond is a thought leader in semantic technology with over 25 years of experience. He is an Enterprise Information Architect at the Bank of Canada. He deploys semantic technology solutions, such as enterprise knowledge graphs, to support AI efforts at the Bank of Canada. He obtained an MA in History from Queen's University and a MIS from UBC.

Corinne Rogers is the Project Coordinator for InterPARES Trust AI. She earned an MAS and PhD from UBC, School of Information, and taught Diplomatics, Digital Records Forensics, and Digital Preservation in the MAS program for 12 years. She has been involved with the InterPARES Project since InterPARES 3 (2007-2012) as a graduate research assistant, and InterPARES Trust (2013-2018) as Project Coordinator.

Appendix B
Dissemination of Results

Members of RP04, *Preserving AI Techniques as Paradata*, participated in numerous presentations, workshops, and panel discussions. Examples of those activities are listed in Tables 1 and 2. Papers written before the publication of this report but not yet published are listed in Table 3.

Table 1. Sampling of articles published in journals and books by the RP04 research study team

Name of author(s) (Lname, fname)	Title	Date (XXXX/MM/DD	Publication	Link to the item, if available
Davet, J., Hamidzadeh, B., Franks, P., Bunn, J.	Tracking the functions of AI as paradata & pursuing archival accountability	2022/06/7-10	Archiving 2022: Expanding Connections Across Digital Cultural Heritage - Final Program and Proceedings (2022): 83-88	https://doi.org/10.2352/issn.2168-3204.2022.19.1.17
	In response to the growing number of experiments in machine learning and artificial intelligence, the InterPARES Trust AI research group proposes a definition of paradata and promotes the consideration of this concept throughout the AI lifecycle. An evaluation of current AI experiments in archives highlights opportunities for paradata-conscious practice.			
Franks, P., Davet, J., Hamidzadeh, B.	Archivist in the machine: paradata for AI-based automation in the archives.	2023/01/20	Archival Science	https://doi.org/10.1007/s10502-023-09408-8
	This paper summarizes the theoretical background of paradata, documentation created during recent experiments with machine learning-based AI, which gestures toward the varieties of paradata already being collected by researchers, and Some of the factors that condition the use of paradata.			
Cameron, S., Franks, P., Hamidzadeh, B.	Positioning Paradata: A conceptual frame for AI processual documentation in archives and recordkeeping contexts	2023/11/16	Journal on Computing and Cultural Heritage, 16 (4)	https://doi.org/10.1145/3594728
	This article reviews the literature on documentation in statistics, the social sciences, and visual heritage; clarifies the relationship between paradata and metadata and XAI; and argues that user-centered explainability can be articulated through paradata, which is positioned as a key tool for addressing AI's challenges to accountability, transparency, and impartiality.			

Cameron, S., Hamidzadeh, B.	Preserving Paradata for Accountability of Semi-Autonomous AI Agents in Dynamic Environments: An Archival Perspective	2024/01	Telematics and Informatics Reports	https://doi.org/10.1016/j.teler.2024.100135
	This paper defines real-time AI systems as time-constrained control systems and argues that they pose accountability challenges that require robust documentation. Using examples from transportation, content targeting, healthcare, and facilities management, it identifies paradata as essential for transparency and for distinguishing human and system agency. Drawing on control theory, it proposes a framework for assessing the documentation and recordkeeping needs of real-time AI systems.			
Henttonen, P., Packalén, S.	Adding Paradata About Records Processes via Information Control Plans	2024/09/17	● In book: Perspectives on Paradata: 215-232 (CC by 4.0), Isto Huvila, Lisa Andersson, Olle Sköld, eds.	https://link.springer.com/chapter/10.1007/978-3-031-53946-6_12
	The authors provide a compelling example of how existing artefacts within administrative processes function as relevant paradata. They aim to stimulate interest in the concept of paradata in recordkeeping and invite discussion of the benefits of understanding certain recordkeeping metadata as paradata.			
Franks, P.	The Crucial Role of Paradata in AI Governance	2024/05/27	SCEaR Newsletter Special Issue 2024: Artificial Intelligence and Documentary Heritage: 54-60. Luciana Duranti and Corinne Rogers, eds. SCEaR (Sub-Committee on Education and Research), UNESCO Memory of the World Programme.	https://unesdoc.unesco.org/ark:/48223/pf0000389844.locale=fr
	AI Governance is necessary to mitigate risks presented by the use of AI tools and technologies. Mitigation requires documentation of the decisions made and actions taken throughout the AI process. The author explores the use of paradata as part of the AI Governance toolkit.			
Cameron, S.	Hybrid Agency and Real-Time Systems: Paradata for Accountability of AI Systems within and beyond Traditional Archives	2024/05/26	SCEaR Newsletter Special Issue 2024: Artificial Intelligence and Documentary Heritage: 61-66, Luciana Duranti and Corinne Rogers, eds. SCEaR (Sub-Committee	https://unesdoc.unesco.org/ark:/48223/pf0000389844.locale=fr

			on Education and Research), UNESCO Memory of the World Programme.	
	Archivists and records managers must be able to identify the risks associated with high-risk AI systems and how paradata can be used to document the decisions made and actions taken throughout the AI process.			
Cohen, M., Jordan, J. M., Sotebeer, M.S., and Stiber, M.	Paradata in Emergency Service Communications	2024/09/17	In book: Perspectives on Paradata: 181-196 (CC by 4.0), Isto Huvila, Lisa Andersson, Olle Sköld, eds.	https://doi.org/10.1007/978-3-031-53946-6_10
	This chapter outlines how iterative, interdisciplinary work to model and simulate emergency services communication systems illuminates the need for a new role for paradata in government archives.			
Huvila, I., Sköld, O., Andersson, L., Friberg, Z., Liu, Y-H.	Paradata: Documenting Data Creation, Curation, and Use	2025/08	Cambridge University Press	https://www.cambridge.org/us/universitypress/subjects/computer-science/computing-and-society/paradata-documenting-data-creation-curation-and-use
	This book covers the basic concepts of paradata and supports practice by covering methods for generating, documenting, identifying, and managing paradata, including formal metadata, narrative descriptions, and qualitative and quantitative backtracking. It also develops a unifying reference model to help readers contextualize the role of paradata within a wider system of knowledge, practices, and processes, and provides a vision for the future of the field.			
Cohen, Fred	Paradata and Forensics in Emerging AI used in Archives	2025/11	All.net	https://all.net/Analyst/2025-11B.pdf
	Abstract: This article summarizes the overall discussion on the use of paradata as a critical component of forensic examination in the context of juridical processes in archives.			

Cameron, S., Franks, P., Huvila, I., Mooradian, N.	"Navigating accountability: the role of paradata in AI documentation and governance." Journal of Documentation	2025/05/12	Journal of Documentation	https://doi.org/10.1108/JD-01-2025-0009
	This paper examines how authoritative documentation requirements address recordkeeping challenges in AI design, development, and use, and how paradata can help meet these challenges. It provides an analysis of the categories of information artefacts required by these different standards.			
Feliciati, P., Duranti, L.	The responsible use of Artificial Intelligence in archives through the use of paradata	2025	JLIS.it Italian Journal of Library, Archives and Information Science,16 (2)	https://jlis.it/index.php/jlis/article/view/636
	This contribution focuses on paradata generated in the context of the use of AI systems or tools in archives to perform archival functions. Its content is primarily the result of research conducted within the context of the international InterPARES Trust AI project.			
Huvila, Isto	Letting AI loose in an archive: Technology to manage or to manage with	2025	Archiv. theori & praxis	https://www.archive.nrw.de/sites/default/files/media/files/ARCHIV-theorie-praxis-online-Heft-1-25.pdf
	Rapidly advancing artificial intelligence (AI) holds great promise to support archival work. At the same time, its diversity complicates grasping its implications through its use in archives and records management tasks and even more so, through its use for generating and processing records in different fields of the society. Understanding AI requires abandoning the simplistic ideas of AI being intelligent like humans and a single monolithic thing.			
Norman Mooradian, Patricia Franks, Amitabh Srivastav	The impact of artificial intelligence on data privacy: a risk management perspective	2025/01/22	Records Management Journal	DOI: 10.1108/RMJ-06-2024-0013
	The purpose of this paper is to increase the artificial intelligence (AI) ethical literacy of information governance professionals by explaining how AI intensifies familiar data privacy issues by virtue of its dependency on data and ability to create new personal information, to explicate emerging privacy enhancing methods and to show their continuity with existing privacy and information governance principles.			
Shinde, G. Kirstein, T., Ghosh, S. and Franks, P.	Tracing the Past, Predicting the Future: : A Systematic Review of AI in Archival Science	2025/10/20	Proceedings of the Association for Information Science and Technology, 62 (1) 659 - 671	https://doi.org/10.1002/pra2.1286

	This review outlines the current state of AI in archival science and records management and lays the groundwork for integrating new techniques to transform archival practices. The research emphasizes the necessity for enhanced interdisciplinary collaboration between AI experts and archival professionals.			
Huvila, Isto	Documenting AI use in humanities research	2026	Proceedings of the Huminfra 2025 conference	https://dspace.ut.ee/items/fc7703e8-30c5-4c93-9102-2572a1b31fa5
	This paper explores the critical need to document the use of Artificial Intelligence (AI) in humanities research and advocates for integrating paradata (process-related meta-information) to capture both technical and human facets of AI use.			

Table 2. Sampling of presentations, professional papers, etc., published by members of the RP04 research study team.

Name of author(s)	**Title of presentation, paper, etc.**	**Date of activity (YYYY/MM/DD)**	**Activity Type (conference, plenary presentation, etc.)**	**Title & Venue - Location if in person (specify if virtual)**	**Link to the item, if available**
Franks, P.	Can AI improve information quality to help support organizational accountability?	2022/05/25	Online Presentation	AIIM True North Webinar	
Jeremy E. Davet, Babak Hamidzadeh, Patricia C. Franks, Jenny Bunn	Tracking the functions of AI as paradata & pursuing archival accountability	2022/06/07	Virtual Conference paper	Archiving 2022: Final Programs and Proceedings, 7-10 June 2022. Society for Imaging Science and Technology, pp 83–88	https://library.imaging.org/archiving/articles/19/1/17
Franks, P., Ghosh, S., Hofman, D., and	Time to Join "AAA": *Are you an AI Aware Archivist?*	2022/07/16	Conference Panel Presentation	2022 NAGARA Conference	https://drive.google.com/drive/u/0/folders/1PkzeYNRP0XFR8SBFO1f9pq1VJLuHWdAW

Mooradian, M.					
Huvila, I., Vats, E., Friber, Z., Börjesson, L., Kaisea, J., Sköld, O.	Automatic identification of archival paradata using artificial intelligence techniques	2022/08/24	Conference Presentation	Final conference of the international 2022-09-08 network Digitization and the Future of Archives: Digital archives, Big Data and Memory, Copenhagen, Denmark	https://www.politics-society.aau.dk/conferences/digi-archives-2022/ (not working)
Franks, P., Hofman, D., Rogers, C., Stiber, M.	*Time to Join 'AAA': Are you an AI-Aware Archivist?*	2022/08/26	Conference Panel - PDF of PPT	*Archives*Records 2022*	https://interparestrustai.org/assets/public/dissemination/SAA_20220826_panel1.pdf
Franks, P.	Capturing and Preserving the AI process as paradata for accountability and audit-trail purposes	2022/10/27	InterPARES Symposium Presentation	1st International Symposium ITrust AI, Lanzarote, Canary Islands	https://interparestrustai.org/assets/public/dissemination/4-Franks_Paradata.pdf
Franks, P. and Cameron, S.	Paradata: Documentation for Responsible Artificial Intelligence	2023/10/17	AIIM Website	Blog Post	https://info.aiim.org/aiim-blog/paradata-documentation-for-responsible-artificial-intelligence
Cameron, S.	Capturing and Preserving the AI Process as Paradata for Accountability and Audit Trail Purposes.	2023/02/20	InterPARES Symposium Presentation	2nd International Symposium, ITrustAI. Abu Dhabi	https://interparestrustai.org/assets/public/dissemination/positioningparadata04-20preprintv21.pdf
Franks, P.	AI and Records Management: Balancing Innovation and Risk	2023/06/23	Online Presentation	SAA Records Management Section Annual Meeting 2023	
Franks, P.	"Paradata: Documenting the AI process for transparency	2023/07/07	InterPARES Symposium Presentation - PDF of PPT	3rd International Symposium, ITrustAI. San	https://interparestrustai.org/assets/public/dissemination/4-Franks_Paradata.pdf

	and accountability			Benedetto del Tronto.	
Stiber, M.	Paradata in Emergency Services Communications Systems	2023/07/07	InterPARES Symposium Presentation - PDF of PPT	3rd International Symposium, ITrustAI. San Benedetto del Tronto.	https://interparestrustai.org/assets/public/dissemination/5-Stiber_ESCS_Paradata.pdf
Franks, P.	In the Pursuit of Archival Accountability: Positioning Paradata as AI Processual Documentation	2023/07/18	Online Presentation and Paper	SAA Research Forum	https://www2.archivists.org/sites/all/files/Franks_In%20the%20Pursuit%20of%20Archival%20Accountability.pdf
Franks, P.	From Data to Disruption: Handling AI Incidents Effectively	2023/10/02	AI Governance Workshop - Day 1 - PPT	IG World 2023 Conference, San DIego, CA	https://docs.google.com/presentation/d/1aVdOEvMMyBpHNc9cq9EiVvMnsMh7DxW3/edit
Franks, P.	Navigating the AI Landscape: From Content Creation through Retention and Disposition	2023/10/03	Conference Presentation - PDF of PPT	IG World 2023 Conference, San Diego, CA	https://drive.usercontent.google.com/download?id=19bje6M2V3luiykb8qEUWA09ulkFz_Hpp&authuser=0&acrobatPromotionSource=gdrive_chrome-list
Richmond, A. and Beauchamp, M.	Architecting accountability in AI: Paradata and AI in commodity markets at the Bank of Canada	2023/10/28	InterPARES Symposium Presentation - PPT	4th International Symposium, ITrustAI, Vancouver, BC	https://interparestrustai.org/assets/public/dissemination/6-Richmond-ArchitectingAccountabilityInterPARESUBCPlenaryandSymposiumPresentation.pptx
Hamidzadeh, B. and Cameron, S.	Paradata, AI, and shared agency and accountability	2023/10/28	InterPARES Symposium Presentation - PPT	4th International Symposium, ITrustAI, Vancouver, BC	https://interparestrustai.org/assets/public/dissemination/7-Hamidzadeh-ITrustAI-Symposium_2023-10-28-2023.pptx
Franks, P.	Paradata: What is it and why do we care?	2024/02/23	InterPARES Symposium Presentation - PDF of PPT	5th International Symposium, ITrustAI, Hawaii	https://interparestrustai.org/assets/public/dissemination/6-FRANKSPARADATATrustAISymposium_pptx.pdf

Franks, P.	Role of RIM in AI Governance	2024/03/13	Webinar	Greater Kansas City Chapter of ARMA (GKCARMA)	
Franks, P.	Converging Worlds: Data, Information, Records	2024/05/09	Virtual Presentation	PRIMO (Pharmaceutical Records & Information Management Org.)- Annual Board meeting	The Role of RIM in AI Governance: Data, Information, Records, and Paradata!
Franks, P.	Demystifying AI: Challenges and Solutions for Records & Information Managers	2024/05/14	Online Spring Conference	Twin Cities ARMA	
Franks, P.		2024/06/11	Virtual Presentation	ACA Virtual Conference 2024	
Cameron, S.	The role of paradata in ensuring trust in AI systems	2025/10/24	Symposium presentation/ webinar	InterPARES Symposium Macerata	
Huvila, I.	Documenting AI use in humanities research	2025/11/13	Conference presentation	Huminfra 2025 conference, Stockholm	

Table 3. Works in progress.

Name of author(s) (Lname, fname)	**Title**	**Date (XXXX/MM/DD**	**Publication**
Franks, P; Mooradian, M.; Fewster, K.; Hodgson, K.	Exploring AI-Enhanced Records Management: A Comparative Analysis of Features in Enterprise Content Management Systems	Submitted 2026/02/26	Records and Information Management Journal (RMJ)

Powell, N.; Franks, P.; Ghosh, V.	Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata	Submitted 3/2/26	Journal of the Association for Information Science and Technology (JASIST)

Appendix C
Reflections in the Words of the Researchers

RP04 researchers were asked to reflect on the team's work since its inception by writing a short (1-2 page) piece that addresses at least two of the four following questions:

1. What can we conclude as a result of the work of the RP04 study?

2. What one change/action does the work produced as a result of the RP04 study suggest individual archival and records management practitioners should be making now?

3. Has/How has the work produced as a result of the RP04 study changed your thinking on how documentation of the AI process can best be accomplished?

4. How has the work produced as a result of the RP04 study changed the way you think we should envision the role of the archivist/records manager for the future?

Rather than including the documents, the comments were extracted and placed under the appropriate questions below. Eight researchers shared their reflections, but not all answered each of the four questions. All researchers agreed to share their names along with their reflections, which are available here. The numbers were added for use in synthesizing and analyzing comments.

Question #1: What can we conclude as a result of the work of the RP04 study?

Researcher #1 **(Isto Huvila):** The key takeaway from RP04 is that paradata is both conceptually and in practice a valuable and necessary concept to complement the documentation of AI systems for diverse purposes, including their accountability, trust, and understandability, as well as their use. It is distinct from parallel concepts such as metadata, explanations (as per XAI), and provenance, both in terms of the information it refers to and as a heuristic device (cf. Huvila et al., 2025) that helps articulate the diverse needs and uses of such information. For records management in general, beyond what InterPARES TRUST AI has been working on (i.e., AI), our study also examines how archives and records management should think about the documentation/keeping of records of the use of technologies other than those labelled as AI. Many different, both technical and cultural, technologies may require similar treatment to that afforded to AI

Researcher #2 **(Patricia C Franks)**: As the use of AI technologies continues to grow, so too does the need for documentation to enhance transparency, accountability, and fairness. Metadata provides structured descriptions to categorize and contextualize datasets, providing details on

provenance, context, and intended use (Bagchi, 2024). At the same time, paradata documents the processes and decisions made during data collection, transformation, and application (Davet et al., 2023; Davet et al., 2022; Huvila et al., 2024). Both metadata and paradata are fundamental components of responsible, ethical AI governance.

Although there is a need to standardize paradata in AI processes, we are not there yet. Organizations must comply with applicable laws, such as the EU AI Act, and consult non-domain-specific guidance documents, such as the NIST AI Risk Management Framework, to determine the documentation required for their AI systems. However, additional guidance must be drawn from domain-specific sources, such as STARD-AI, which assesses diagnostic accuracy in healthcare (Sounderajah et al., 2025).

Researcher #4 (**Fred Cohen)**: Paradata becomes the last hope of meeting the requirements of surety for the archives by identifying the information required to meet otherwise unmet objectives and the methods by which that information may be used to achieve the goals. The set of information and methods may then be provided as paradata. As this is done, the paradata is also codified in archival records. In essence, paradata may be a set of documents associated with the archives (and perhaps contained within them) required to meet juridical requirements not otherwise met by the archives.

Researcher #5 **(Norman Mooradian):** Cameron et al. explore the role of paradata in records creation in support of AI transparency. In particular, they investigate how authoritative documents such as laws, model laws, and governance models explicitly and implicitly enjoin or promote the creation of records in the development and deployment of AI systems and subsystems. They also explore how the concept of 'paradata' relates to the more established concepts of records and documentation, and how paradata helps us interpret recordkeeping requirements in relation to AI and information and data relevant to creating records.

A central problem for records and archives is how to define, identify, and create an "AI record." (Mooradian, 2019) An AI record is a record of decisions and/or actions taken by an AI-enabled system relative to some interest or right of individuals and groups. Because so many AI processes are opaque (black box) to their operators, capturing records of the reasons or bases of decisions is especially problematic. National and transnational laws and standards attempt to promote transparency in various ways. Some of the authoritative documents examined by Cameron et al. (2025) included instances where recordkeeping of some sort was explicitly required, though the injunctions often lacked specificity. More often than not, however, they found that the authoritative documents required the creation of documentation, the collection of information, or the undertaking of activities that implied a requirement for the creation of specific records. Moreover, Cameron et al. found that the concept of paradata provided a unifying vision of many such implied records. In particular, the concept of paradata provided a conceptual tool for identifying certain types of data and information created during the AI lifecycle that are essential to building adequate records. For records and archives, the implication is that there is an opportunity to work with stakeholders across the AI lifecycle to identify paradata available and/or required for the construction of an AI record.

Researcher #6 **(Rae Lynn Haliday):** As a result of the RP04 study, I now understand that paradata produced by the digitization of legacy records with the assistance of Artificial Intelligence (AI) can help educate information professionals about where data quality and quality control problems can arise. Furthermore, it can facilitate business process improvements and provide a path to long-term protection, preservation, and integrity of records in digital legacy systems. Using AI to assist in the digitization, organization, and indexing of legacy records into modern technology systems requires effective paradata and related analysis for 1) transparency and user trust, 2) identification of errors and biases, and 3) flagging of substandard image or AI quality for review, as those records may be indexed with lower accuracy.

This is critical as we've witnessed scenarios like microfilm quality metrics, including image quality, handwritten records that are hard to read or have errors, have been digitized by multiple systems, and have lower Optical Character Recognition (OCR) scores. Using AI to assist users in indexing digitized microfilm of handwritten records into modern systems may require more human ingenuity and manual resources than initially anticipated at the beginning of this study. To optimize the process, we must now identify the steps that most impede digitization and require additional attention and resources.

Researcher #8 **(Alex Richmond)**: Here are some thoughts on what I'd call applied learnings around implementing paradata in real institutional contexts:

Key lessons that surfaced:

- Capturing paradata manually is expensive — automation is essential
- Documentation must integrate into workflows from inception; retrofitting is inefficient.
- Governance teams and model developers often use different terminology; shared definitions are critical.
- Paradata is only useful if structured and linked to lifecycle artifacts.
- There is a need to balance high-level guidance with domain-specific detail in checklists and taxonomies.

Collectively, these findings represent a pragmatic awareness that *successful paradata systems require cultural adoption as much as technical design.*

Question #2: What one change/action does the work produced as a result of the RP04 study suggest individual archival and records management practitioners should be making now?

Researcher #1 **(Isto Huvila):** The significant action every individual archival and records management practitioner should take now is to understand the importance of keeping information (or paradata) on AI systems and how paradata is used in processes that should generate records. Another parallel observation, consistent with the previous findings from the InterPARES project cycle, is that the notion of record is both changing and persisting. This is also evident in the case of paradata, which takes many forms and depends heavily on the technologies involved in the processes it is intended to document.

Researcher #2**(Patricia C Franks)**: A good first step is for archivists and records managers to understand the AI process as it relates to the systems they employ (or could employ) and the steps that must be taken to document that AI process, including the use of both metadata and paradata.

Researcher #3 **(Scott Cameron):** A significant challenge in ensuring adequate paradata is the education of archivists. Archivists should ask what information a tool provides when they are applying it, both before use and throughout its use. We need system-level explanations of how tools work, and we also need information about particular decisions made by the tool. This all constitutes paradata. Beyond the increased transparency this can encourage, the concept/framework is fruitful for building AI literacy within the information professions. If we can get archivists to ask about available paradata for each phase of the AI lifecycle, they will already have adopted a toolkit of critical analytical tools, enabling a deeper understanding of AI tools.

Given the limitations of resources, time, and AI skills in archival institutions and among their workers, AI tools are as or more likely to be off-the-shelf systems implemented on a SaaS basis than to be implementations (as is or customized) of more transparent, open-source models. Accordingly, developing the AI literacy necessary to interrogate and compare the paradata offered by competing AI system vendors is, in my view, a crucial competency for the procurement of AI tools for use in archives. Archivists will not be able to shape the nature of AI for archives if we do not understand it; a first step here is to gain the capacity to analyze the strengths and shortcomings of existing AI tools imported for archival use. Beyond efficacy in performing the primary task, offering useful paradata should be a core assessment factor of any AI tool – and archivists will need to be able to assess this paradata accordingly.

Researcher #4 (**Fred Cohen)**: The design of paradata requirements for archives should presumably be driven by the Juridical requirements of the archives. It is obvious that this includes the legal processes the archives support and, as public records, the requirements of the relevant agencies and other sources of information the archives collect, preserve, and make available for use. Defining the uses and requirements of the sources should usually be part of the design of the archival processes, and this will generally be reflected in the applicable elements of the COP (Chain of Preservation) table selected for use, the descriptions of the archives and their functions, and the laws and regulations governing them.

These elements must be translated into the specific activities of the archives and linked to the mechanisms that provide them, as well as to the level of certainty associated with the methods applied. If and to the extent the methods provide adequate certainty that the objectives they support will be accomplished under the design basis threat (which presumably has to be defined based on the same information), additional requirements for attaining unmet or inadequately met protection objectives for the desired level of certainty and that cannot reasonably be provided based on the other available information in the archives, paradata is the 'bucket' for the 'everything else' required.

Researcher #5 **(Norman Mooradian):** For records managers and archivists, Cameron et al. provide several insights that should be useful for constructing an AI record. One insight is that compliance mandates do not always state their recordkeeping requirements explicitly or provide specific guidance. They do, however, explicitly require activities such as testing, documenting, communicating, and training that carry with them information requirements and data trails. These various forms of information and data all fall under the concept of paradata, and they likely are elements of an AI record. Records managers and archivists should therefore be sensitive to, and devote effort to identifying, language in authoritative writings that implies that such paradata exists or should exist.

The second insight provided by Cameron et al. is that some paradata is human-made, while other paradata is machine-generated. Both human-made and machine-generated paradata may be intended for use as part of a record, but more often than not, they are intended for a purpose earlier in the AI lifecycle. Further, human-made paradata in the form of policies, communications, and documentation is often best suited for understanding AI development and deployment globally or generally. Machine-generated paradata, by contrast, is often more useful as part of a record of specific operations of AI systems, e.g., decision-making. Records managers and archivists should therefore be aware of these different types of paradata and understand their value in creating specific types of AI records.

The takeaway for records and archives is that AI governance will increasingly require activities that generate significant amounts of paradata of varying kinds. This paradata will be critical to the creation of AI records, which are inherently challenging because of the opacity of certain forms of AI. Records and archives will therefore need to broaden their view of what constitutes a record and where such information can be found. Paradata will serve as a valuable concept for such an expansion.

Researcher #7 **(Jeremy Davet):** The work by the RP04 study within the InterPARES Trust AI research group over the past handful of years has, if nothing else, highlighted the need for greater urgency on the part of all archivists and record management practitioners; both to assess the AI-based tools which they are now or may soon rely, and to capture critical information about their effects on record bodies. While robust documentation describing the function and capabilities of commercial software and automated archival functions has proven to be an asset to recordkeepers since the early days of the digital age, the proliferation of AIs demanded the definition of paradata.

The maturation of large language models (LLMs) during the period of research has only served to reemphasize the need for collaborative urgency and action, as human-information interaction within and without the archives has begun to evolve along unexpected avenues. To forestall the disappearance of fact and meaning into the miasma of model hallucination, the work of RP04 suggests that individual recordkeepers should use the framework of paradata to interrogate the AI tools that they and their parent organizations are being sold.

Researcher #8 **(Alex Richmond):** This post addresses challenges, open questions, and research gaps. I wanted to also highlight some persistent challenges:

- **Definition harmonization**: Different disciplines interpret paradata differently (archives, ML engineers, governance).
- **Maintaining paradata quality** across evolving datasets and models.
- **Balancing transparency with security** when documenting sensitive AI systems.
- **Deciding responsibility**: Who records paradata, and who verifies its accuracy?
- **Capturing paradata in real-time or high-velocity AI systems**, including semi-autonomous agents.

Addressing these questions is essential for making paradata sustainable and operationally viable.

Question #3: Has/How has the work produced as a result of the RP04 study changed your thinking on how documentation of the AI process can best be accomplished?

Researcher #1 **(Isto Huvila):** Even if I knew better from the start, my hunch was that we might end up with a nice and tidy paradata standard that would provide a starting point for all relevant paradata to document an AI system. Now I don't mean that standards would not be needed. Still, I am reasonably sure that the problem cannot be solved by standards alone; instead, it requires thorough work within individual domains and systems to understand how basic archival concepts, such as accountability, function within the context of AI-powered processes. We are not really there, but work as Scott and Babak did with autonomous systems and as you, Norm, Scott, and I did to identify what current efforts to support AI transparency comprise and lack, are essential steps in this direction. Perhaps the worst that could happen now is that a technical working group would quickly try to develop a seemingly generic standard for AI paradata based on, e.g., the InterPARES definition of paradata (or any other) without taking time to consider not only what things might be relevant to document but also what AI use looks like in specific contexts.

What I find critical when considering how to document AI process is to consider how formally structured paradata complements and is complemented by diverse records relating to AI process that go parallel to, e.g., the administrative process within which AI is used. In this sense, paradata can easily consist of records and form an 'archive' of its own. Another aspect to consider is also what is meant by AI process and to what extent they are similar to and different from other processes. At the beginning of RP04, I perhaps thought that the whole thing was about documenting a technology or tool and its use. But now – even if I am critical of anthropomorphizing AI – I am probably more inclined to think that the AI process has probably much more in common with sociotechnical non-AI processes than I first thought. The AI itself might not be as important in the end as how it was trained on data and how its outputs were assessed and acted upon, much as we have traditionally provided records of people as data processors and decision-makers.

Researcher #2 **(Patricia C Franks)**: Manual documentation is not practical. The creation and collection of paradata must be embedded within AI systems where possible to document procedural steps in real time, thereby making AI models traceable and auditable throughout their lifecycle (Longpre et al., 2024). A better understanding of how metadata and paradata can be integrated into the AI workflow is needed.

While paradata generated through the AI system (workflow) may be captured automatically, some paradata will exist in other forms for different purposes, such as impact statements, AI policies, employee training records, and more. There needs to be a way to aggregate all supporting documentation regardless of where it resides. This calls for a portal that contains an inventory of all AI projects within the organization, the required documentation (for various stakeholders), and links to those resources (e.g., datasheets, model cards, vendor documentation).

Researcher #3 **(Scott Cameron):** Earlier on in the study, I had a bias towards technical approaches to paradata. I had assumed that the biggest obstacle to providing relevant paradata was in AI system design: AI tools would simply not provide sufficient technical logs about their ongoing operation to offer useful paradata sufficient to hold users accountable in a precise manner. I no longer think this is the case. Automated technical logs of AI outputs are, in some cases, key paradata artefacts. Still, I tend to understand human decisions as more prominent factors in relevant paradata than technical data. Questions like those listed below are organizational rather than technical:

- Who decided on the AI tool to use in this case?
- How was the tool adapted for the particular use case?
- What quality assurance measures were performed before implementation?
- What ongoing monitoring was conducted?
- How is the organization notifying relevant data subjects or other interested parties of the AI use and its possible consequences?

These questions, I would argue, are as relevant or more relevant than any system-level output logs offered by an AI tool. While the technical black box may be a significant obstacle to AI applications, the most pressing questions that must be clearly answered in high-risk applications, as ever, remain organizational and human.

Researcher #4 **(Fred Cohen)**: Paradata might be well described as 'everything else.' That is, we have the data associated with records, we have the metadata like filenames, location in the fonds, inherent properties of the digital forms the records take, protection settings, authorship information, date and time stamps, and the rest of the elements of the Chain of Preservation (COP) table, and yet, even with all of this information, we may find it hard to do the work required for examination. As a general rule, absent any other constraints, it is a good idea to preserve the software and mechanisms, data, and factual content required to reproduce the necessary functions of the archives. Of course, there are always other constraints present. Given those constraints, several additional approaches can be employed to realize many of the cost and performance advantages of emerging AI while mitigating the uncertainties inherent in their use through appropriately generated and selected paradata.

Question #4: How has the work produced as a result of the RP04 study changed the way you think we should envision the role of the archivist/records manager for the future?

Researcher #1 **(Isto Huvila):** As it is quite apparent that records keep changing their shape, an archivist/records manager of the future is hardly the stereotypic character that managed records (as things) but rather should be a specialist of documentary accountability and how it is implemented on the levels of technical and sociotechnical systems through producing records of different kinds that fulfil the specific function in individual cases.

Researcher #2 **(Patricia C Franks)**: As was the case throughout time, archivists and records managers must become familiar with the new technology and terms introduced into the workplace. Just as information professionals in the library field led the development of metadata standards, archivists and records managers can play a vital role in developing paradata standards. Archivists and records managers must join interdisciplinary teams that evaluate the role of metadata and paradata in supporting compliance with evolving AI regulations and mitigating bias in algorithmic decision-making.

Researcher #3 **(Scott Cameron):** There is an opportunity here for archivists to model responsible and transparent AI implementations for other professions and for society at large. Demonstrating a grounded and critical attitude towards AI implementation is not difficult, but few individuals or organizations avoid succumbing to the hype surrounding new technologies. The technology has significant limitations that should be transparently documented and conveyed. Archivists can and should model this grounded approach to technology by providing useful, clear paradata to archive users.

Researcher #6 **(Rae Lynn Haliday):** AI affects the work of modern information professionals and underscores the necessity of capturing and documenting AI paradata as a standard part of managing digital records and systems, no different from metadata requirements and related standards, but equally important. AI-driven processes (e.g., OCR, HTR, and record linking) create records in ways that are not always transparent and can evolve rapidly. AI paradata helps standardize the use and application of the technology. Making it an essential component of technology standards will ensure accountability, interoperability, preservation, and ethical responsibility.

Researcher #7 **(Jeremy Davet):** The work produced as a result of the RP04 study has convinced me that the roles of archivist and records manager must both expand and contract, in some respects.

As regards contraction, what I really mean is that the roles must more acutely emphasize the development, preservation, and transmission of technological literacies, and emphasize the management of automated systems. While we now have the intellectual framework to speak productively about paradata, archivists and records managers must now work to operationalize paradata, to refine it from the functioning of automated systems, and to make collective meaning from it. Professionally established and developing recordkeepers alike must become more conversant in the language of computer science and digital infrastructure if they are going to advocate effectively for archival principles in an increasingly automated world. In some

cases, this may require foregoing the development of other knowledge and of intellectual generalism.

The roles of archivist and record manager should also expand. The definition of paradata allows the recordkeeper to build bridges between disciplines—articulating the needs of the record system to information professionals designing AI, refining technical subjects for end users of varying levels of technological literacy, describing the effects of increasing automation on information-seeking behavior to academics, and so on. The work of the RP04 study demonstrates how fruitful these kinds of interdisciplinary collaborations can be, and models (in miniature) how they might be constituted. The archivist or record keeper can no longer be contented in limiting their efforts to the archives, or their peers to the like-minded.

Appendix D
Use Cases – Supplemental Information

Use Case #1 - Use Case for Automated Data Recognition & Extraction for Indexing Digitized Images of Animal Records from Legacy Microfilm

Submitted by Rae Lynn Haliday, MBA, CRM, CIGO, FAI, the Curator of Animal Management Services—Animal Records, Transport and Compliance at the St. Louis Zoo.

The Saint Louis Zoo's mission is to conserve animals and their habitats through animal management, research, recreation, and educational programs that encourage the support and enrich the experience of the public (https://stlzoo.org/about/mission). This use case is relevant to 40 years of legacy records (created from 1969-2009) that were microfilmed and that are currently being digitized. These records represent four separate record series and are classified as permanent retention on the zoo-specific records retention schedule. The original camera masters are stored in an off-site microfilm vault; as a local county government, the Zoo qualifies for specific resources including microfilm storage and digitization services.

The Saint Louis Zoo's Animal Management Services (AMS) department is working with the State of Missouri Local Records program to digitize the legacy microfilm. Through its virtual and hybrid internship program, AMS hosts interns from Master's in Library and Information Science (MLIS), RIM Graduate Certificate, or other higher education programs for Records and Information Management (RIM), to assist with the project. The goal of the project is to digitize these records so that this information with enduring reference, research and legal value can be accessed online internally by authorized staff and the microfilm reader printers and duplicate microfiche can be dispositioned.

Testing was conducted in conjunction with the vendor to evaluate features that are part of the current Electronic Document Management System (EDMS) in use for workflows, active, inactive and legacy animal records, to determine if Artificial Intelligence (AI) applications could be used to speed the work processes, including file splitting (extracting) and learning document types for indexing.

Intelligent Indexing Research for Application in the Digitization Use Case

1. 10-20 sample Animal Transaction Files were tested. Among the challenges identified was that the Intelligent Indexing feature had difficulties reading or interpret handwritten information due to image quality and poor handwriting.
2. Splitting files (extracting) individual file information was complicated by the feature's difficulty in recognizing the record series number in terms of when one file ended and

the next one began. This was the case with typed information as well as handwritten. Even with a barcoded label, the quality of the image was cited as an issue by the vendor.

3. Image quality of the characters in the record series number, and the background interference inherent with the migration of paper that was microfilmed and then digitized (digitized film to PDF image), presented as the primary limitations to using the Intelligent Indexing feature for this project**.** These Images were filmed at ANSI Standards with 200 ppi/DPI or 300 ppi/DPI for documents with small fonts, signatures, or those intended for optical character recognition.
4. The Intelligent indexing feature was adding the screenshot of the file folder with the next file to the current file during the splitting (extracting) process and/or it would continue adding many documents from other files (i.e. a file that had 20 pages now had 100 spanning across other files).

   This would have required staff to breakout each file after splitting process to perform quality control to address these issues.

5. Two approaches were used to split (extract) files digitized from the original microfilm rolls: 1) Adobe Professional was used initially to select, organize and extract all documents of each individual file (record series) that was digitized and then upload and index in DocuWare 2) Upload the single PDF/A of the digitized microfilm roll to document tray in DocuWare, and unstaple and index documents in PDF/A.

   The latter approach required less downloading and reuploading of documents to additional systems and provided a direct path for delivering the documents from staff to interns for holding and indexing information into the targeted system. It was determined that Intelligent Indexing would probably not work to speed the digitization process for the majority of these legacy files.

   However, we did discover that the Intelligent Indexing feature could be used: 1) When capturing data from different forms that aren't conductive to using templates, look for the key word and capture everything around it. 2) When working off of first-generation paper for best results. 3) When processing PDFs received via email.

Bar Code and Forms Module to Split Files

1. The Bar Code and Forms Module may be a stronger splitting feature. First generation bar codes work best. If the bar code is second or third generation, the quality of splitting will be impaired, as we saw when the sample files were tested.

AI Tool - Intelligent Documenting Processing (IDP)

IDP uses AI to increase the reading accuracy of images and characters to improve data capture. It offers Classification of documents and Automation of the records management processes.

IDP recognizes the layout of documents for consistency in application, but it doesn't retain the data. It scrubs data and retains only the document layout.

IDP was considered given the amount of handwritten information in the legacy files that we have to digitize. At this time, however, we chose to use the features in hand that don't require additional costs due to budget constraints.

Summary

To summarize, after the evaluation and testing was conducted to determine which AI tools, if any, were available to automate and speed the process of indexing the digitized images of Animal Transaction Files previously held on microfilm, we learned more about the capabilities of the system we already have access to without the need to incur unnecessary expenses for this project.

We better understand how to evaluate the risks of using AI. For example, the risks for using features like Intelligent Indexing or Intelligent Document Processing can change depending on whether you are using locally-hosted or cloud-hosted services. We received the AI White Paper from our vendor, and it explains the types of AI and the benefits of each, and which ones are used or not used. It explains the models and how AI is used to the customer's advantage across its various products. It further details how security of information and anonymization of the data is handled as well as cross-customer modules. As a result of this project, we now have a better understanding of AI's use in records management, its risks, pros and cons, limitations, and its ability to work with different media (i.e., paper, microfilm, paper that has been microfilmed and then digitized).

Ultimately, it was determined that the current process being utilized for the digitization, inventory, indexing and quality control for records identified for this use case is the best path forward due to the human component required for handwritten information, experience with the unique records series, and for budgetary constraints. While the records span 40 years, the volume is manageable. State resources are providing the means for digitization of the legacy microfilm records. AMS staff is reviewing and uploading the information to document trays and training interns on the indexing and quality control procedures. Interns from higher education, records centered-degree programs are learning about and obtaining real-world experience in modern tools for records preservation, and how to inventory, index and analyze legacy records for possible inclusion into the globally-recognized ECM.

The information that was gained from the evaluation and testing of modern AI tools can be used moving forward to enhance automation, capture, classification, OCR quality assessments, metadata validation and general quality control for born digital records.

https://start.docuware.com/blog/document-management/docuware-ai-hub?utm_campaign=29506614-homepage-update-25&utm_source=en-us-homepage&utm_medium=cta&utm_content=ai-hub

USE CASE 1: APPENDIX A

Background and Vendor Process for Evaluating the use of AI for Intelligent Indexing of Digitized Documents

1. Review sample data (documents and/or electronic image files) 10 to 20 samples
2. Determine metadata required to recognize/extract
3. Review complexity of document/image layout
    a. Document Type: 1 or multiple document types?
    b. Document Layout: repeating layout or random layout?
    c. Complexity of Data: simple (2 to 3 fields) or many index values, uniqueness of index fields, or line item/table detail?
    d. Consistent or varying document Page Lengths: how will the images be split into separate documents? What is the trigger for the splitting?

Addition Information

- Monthly volume of images/documents
- Physical condition of paper documents to be scanned
- Paper Documents to be scanned or Images files (or both)
- Format of electronic image files (PDF, TIFF. JPG, etc.).
- Source of electronic files (email attachments, file folders, etc.)

Configuration of Use Case for Testing & Deployment

This will determine the destination for the classified and indexed documents/file, and how they will be searched for and retrieved. Also, any workflow process will be determined and processes for exception handling will be established.

1. Determine (or create) the repository file cabinet.
2. Determine (or create) the index fields for the files and determine which will be filled with the extracted data.
3. Create DocuWare Input Tray for the images (The Intelligent Indexing processing will be configured to begin when documents/images are input into this tray.
4. Determine sources for the documents/images (scanners, email in-baskets, drop-and-drop from folders or drives, etc.)
5. Determine trigger data on documents/images to be used for splitting files into separate documents
6. Select document/files test batch
7. Input sample document/files sets into the input tray for processing
8. Review results: Splitting, Index value capture using the Red, Yellow, Green indications
9. Train Red or Yellow, verify correct data capture, rope around index value to indicate correct index value location and train system. Green, results are accurate.
10. Re-import a second test batch

11. Review/train results
12. Re-import a third batch (if needed)
13. Review/train results

After the system goes live, import results will be monitored and any errors that are identified (Red or Yellow indicators) can be trained by the user.

Storing of files/images of Animal Transaction Records into DocuWare digitized from records previously archived on Microfilm reels and rolls using Intelligent Index capture and enhance indexing via Auto-Index.

Review sample data (documents and/or electronic image files)

Files are presented to the Animal Management Service's office in PDF format that have been digitized from Micro-Film reels and rolls. Each file is a grouping of a single reel or roll per file. The PDF file naming convention used is XXX.pdf where the XXX represents the reel or roll number.

Reel and roll files from the 1970s, 1980s, and 1990s are primarily identified with hand-written 'cover' pages for each document contained on the reel/roll. This represents approximately 2,250 documents. Files from 2000s to present are identified with machine printed cover page. Each subset (70s to 90s and 2000 to present) each have a consistent cover page layout.

Appendix D
Use Cases – Supplemental Information

Use Case #2: The Role of Paradata in the Bank's LLM Journey

Submitted by Alex Richmond, Bank of Canada

The Bank of Canada is utilizing paradata to refine internal models, such as the **Oil Market LLM** and the **Business Outlook Survey (BOS)** analysis tool. Key insights from these projects include:

- **Systematic Labeling:** Understanding what drives a model to apply certain labels helps identify and correct learning errors.
- **Content Evaluation:** Shifting from raw counts to the **share of paragraphs** a theme appears in provides a more accurate assessment of content.
- **Rationale:** Adding rationales to model outputs helps manage normalization challenges and provides deeper context for positive or negative ratios.

Experimental Paradata Framework

The Bank has developed an experimental framework to standardize the capture of metadata across the entire **Machine Learning (ML) lifecycle**. This framework is categorized into several core areas:

- **Core Descriptive Metadata:** Includes unique identifiers, versioning, creators, and creation dates.
- **Training Data Paradata:** Tracks specific data sources, sizes, date ranges, licensing, and detailed preprocessing steps like cleaning and tokenization.
- **Architecture and Training:** Documents model types (e.g., Transformers), hyperparameters, hardware (GPUs/TPUs), and software frameworks used.
- **Evaluation and Ethical Considerations:** Records performance metrics (F1-score, accuracy), benchmarking against other models, bias mitigation efforts, and known limitations.

Automation and Prompt Engineering

To reduce manual burden and ensure consistency, the Bank emphasizes **automating paradata collection** by integrating with tools like **Git, MLflow, and DVC**. Furthermore, **Prompt Engineering** is treated with the same rigor as code; prompts are version-controlled and associated with metadata regarding their intent, task, constraints, and target model.

Distinction from Standard Model Cards

While *Hugging Face Model Cards* provide a high-level overview, the Bank's **Paradata Framework** is more comprehensive. It explicitly addresses **data governance**, offers a **standardized schema** for interoperability, and focuses on the **automated preservation** of all artifacts generated during the ML workflow.

Appendix D
Use Cases – Supplemental Information

Use Case #3: Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata

Submitted by Nancy Powell, SJSU GRA; Patricia C. Franks, SJSU Professor Emerita; and Souvick Ghosh, SJSU Assistant Professor

Rather than include documentary evidence in the Final Report, a paper was developed based on research conducted between September 2024 and February 2026. The title, "*Governing Clinical AI as Information Infrastructure: A Lifecycle Documentation Architecture of Metadata, Provenance, and Paradata" will expand upon the information provided in this report.*

The paper is in its final stages and will be submitted to the Journal of the Association for Information Science and Technology (JASIST).

https://asistdl.onlinelibrary.wiley.com/journal/23301643

Appendix D
Use Cases – Supplemental Information

Use Case #4: InterPARES CU05 Project: Information Framework for Documenting Case Studies

Submitted by the CU05 team: Maria Guercio, Stefano Allegrezza, Francesca Magnoni, Grandi Massimiliano, Bruna La Sorda, Maria Mata Caravaca

The CU05 team has developed a template to track paradata to document case studies based on the AI application lifecycle and project phases of InterPARES Trust AI, which was completed January 5, 2025.

Since then, the approach has been applied to the first case study and modifications to the template have resulted. Details on this project will be provided in the CU05 Final Report available through the InterPARES Trust AI website, Dissemination tab, at: https://interparestrustai.org/trust/research_dissemination

Appendix E

AI & Paradata topics for Inclusion in Archives & Records Management Training and Education

The topics recommended for inclusion in Archives & Records Management Training and Education are based on the belief that archivists' and records managers' future role is clearly shifting toward AI-aware, system-centered, accountability-focused recordkeeping. Many of these topics may already be a part of formal education and training programs, but the list can also be used by information professionals seeking to identify and fill gaps in their current knowledge and skills.

Mod 1: I Literacy for Recordkeepers (It's not all fun!)

- What AI is (and isn't)
- Machine learning basics, LLM basics
- Why AI outputs are probabilistic
- Common failure modes (hallucination, drift, bias)

Mod 2: Paradata: Theory → Practice

- What paradata is and why it matters
- Paradata capture workflows
- Paradata as "evidence of record creation"
- How paradata differs from metadata, provenance, and XAI

Mod 3: Standards Development for Paradata

- How standards are made
- What a generic paradata standard might include
- Crosswalks: metadata ↔ paradata ↔ XAI
- Interoperability and documentation models
- Consideration of Domain-Specific Paradata requirements

Mod 4: Records as Sociotechnical Systems

- Moving beyond "records as objects"
- Systems thinking applied to recordkeeping
- How organizations, software, policy, and humans co-produce records

Mod 5: Documentary Accountability & Evidence in Automated Systems

- Accountability as a recordkeeping function
- What counts as evidence when AI is involved
- Verifiability, traceability, and explainability

Mod 6: Algorithmic Bias, Fairness, and Documentation

- Bias in training data and workflows
- What recordkeepers can document and control

- Risk documentation as a recordkeeping responsibility

Mod 7: AI Governance, Law, and Compliance (Applied)

- AI regulation basics (high-level)
- Compliance documentation strategies
- Governance models and audit readiness

Mod 8: Technical Infrastructure for Digital Recordkeeping

- APIs, databases, cloud platforms
- Data pipelines and automated processing chains (OCR/HTR/linking)
- Versioning and system evolution

Mod 9: Audit Trails, Logging, and Transparency

- What logs matter and how to preserve them
- System audit trails as archival records
- Documentation of model versions, parameters, and training sets (when possible)

Mod 10: Preservation of AI-Generated and AI-Processed Records

- Preservation risks in AI-mediated records
- Capturing context and process
- Long-term interpretability (future users' understanding of how a record was made)

Mod 11: Communication & Translation Skills

- Explaining technical processes to non-technical users
- Translating archival principles to engineers and system designers
- User-facing transparency (what end users need to know)

Mod 12: Interdisciplinary Collaboration & Team-Based Work

- Working in mixed teams (archives + IT + legal + compliance + data science)
- Requirements gathering and documentation practices
- Participating in system design and procurement

Mod 13: Critical Tech Mindset (Anti-Hype Training)

- How to evaluate new tools responsibly
- Documenting limitations and uncertainty
- Avoiding “innovation theater”

Mod 14: Professional Identity: The “Expanded and Contracted” Role

- The shift toward deeper technical specialization
- What may need to be deprioritized
- Building career pathways: archivist-as-governance-specialist