JOURNAL OF SUSTAINABLE COMPETITIVE INTELLIGENCE

DOI: https://doi.org/10.37497/eagleSustainable.v17i.1352

ARTICLE

A Pre-Deployment Competitive Intelligence Protocol for Assessing Multimodal Cultural Heritage Repositories

Protocolo de Inteligência Competitiva Pré-Implementação para Avaliação de Repositórios Multimodais de Patrimônio Cultural

Qiuying Li

School of English and Advanced Translation, Beijing Language and Culture University, Beijing, China.

ABSTRACT | Purpose: This study tests a pre-deployment Competitive Intelligence protocol for converting repository-level evidence into clearly bounded management decisions without implying that the framework has already been implemented or validated within an organization. **Methodology/Approach:** The study examined 35 files deposited in Harvard Dataverse and one directly accessible README file. The analytical procedures included repository inventory, modality and storage profiling, TF-IDF analysis, heritage-theme coding, calculation of a five-component Repository Processing Readiness Index, sensitivity testing under an alternative weighting specification, and a role-based decision simulation. No external users or institutional decision-makers participated. **Key Findings:** Processing-readiness scores were 90.0 for text, 40.2 for video files, 40.0 for compressed archives, and 24.4 for model files. Under the alternative weighting specification, compressed archives ranked ahead of videos. The resulting decision rules supported an archive-extraction pilot, later-stage video processing, stronger documentation controls, and the temporary deferral of model weights. Cultural value at the asset level, organizational utilization, stakeholder feedback, and institutional outcomes were not tested. **Originality/Relevance:** The study provides an empirical Competitive Intelligence output at the repository level while maintaining a clear distinction between processing readiness, cultural significance, and organizational use. This distinction reduces information asymmetry and prevents technical accessibility from being interpreted as evidence of cultural or strategic value. **Theoretical/Methodological Contributions:** The proposed protocol connects Competitive Intelligence planning, environmental scanning, intelligence governance, information asymmetry, evidence grading, decision rules, sensitivity analysis, and validation thresholds within a transparent pre-deployment framework for multimodal cultural heritage repositories.

Keywords | Competitive Intelligence. Pre-deployment protocol. Repository processing readiness. Intelligence governance. Multimodal cultural heritage. Decision simulation.

Qualis B1 CAPES 2021-2024

Indexed by Scopus

ELSEVIER

Received: 12 Jun. 2026

Revised: 21 Aug. 2026

Accepted: 27 Aug. 2026

e-ISSN: 3085-7813

Corresponding Author: Qiuying Li – E-mail: louiselee@pku.org.cn

Editor in chief: Altieres de Oliveira Silva – Alumni.In Editors

How to cite this article: Li, Q. (2027). A Pre-Deployment Competitive Intelligence Protocol for Assessing Multimodal Cultural Heritage Repositories. *Journal of Sustainable Competitive Intelligence*, 17, e01352. https://doi.org/10.37497/eagleSustainable.v17i.1352

RESUMO | Objetivo: Este estudo testa um protocolo de Inteligência Competitiva em fase de pré-implementação para converter evidências em nível de repositório em decisões gerenciais claramente delimitadas, sem sugerir que o framework já tenha sido implementado ou validado em uma organização. **Metodologia/Abordagem:** O estudo examinou 35 arquivos depositados no Harvard Dataverse e um arquivo README diretamente acessível. Os procedimentos analíticos incluíram inventário do repositório, caracterização das modalidades e do armazenamento, análise TF-IDF, codificação de temas relacionados ao patrimônio cultural, cálculo de um Índice de Prontidão para Processamento do Repositório (*Repository Processing Readiness Index*) composto por cinco componentes, teste de sensibilidade com uma especificação alternativa de ponderação e uma simulação de decisão baseada em papéis organizacionais. Não houve participação de usuários externos ou de tomadores de decisão institucionais. **Principais Resultados:** Os escores de prontidão para processamento foram de 90,0 para textos, 40,2 para arquivos de vídeo, 40,0 para arquivos compactados e 24,4 para arquivos de modelos. Sob a especificação alternativa de ponderação, os arquivos compactados ficaram à frente dos vídeos. As regras de decisão resultantes sustentaram a realização de um projeto-piloto para extração de arquivos compactados, o processamento de vídeos em uma etapa posterior, o fortalecimento dos controles de documentação e o adiamento temporário dos pesos dos modelos. O valor cultural em nível dos ativos, a utilização organizacional, o feedback dos stakeholders e os resultados institucionais não foram testados. **Originalidade/Relevância:** O estudo fornece um resultado empírico de Inteligência Competitiva em nível de repositório, mantendo uma distinção clara entre prontidão para processamento, relevância cultural e utilização organizacional. Essa distinção reduz a assimetria informacional e evita que a acessibilidade técnica seja interpretada como evidência de valor cultural ou estratégico. **Contribuições Teóricas/Metodológicas:** O protocolo proposto conecta o planejamento da Inteligência Competitiva, a varredura ambiental (*environmental scanning*), a governança da inteligência, a assimetria informacional, a classificação das evidências, as regras de decisão, a análise de sensibilidade e os limiares de validação em um framework transparente de pré-implementação para repositórios multimodais de patrimônio cultural.

PALAVRAS-CHAVE | Inteligência Competitiva. Protocolo de pré-implementação. Prontidão para processamento de repositório. Governança da inteligência. Patrimônio cultural multimodal. Simulação de decisão.

1 INTRODUCTION

Digital heritage collections now combine images, descriptions, catalogue records, geospatial labels, conservation notes, institutional metadata, videos, and model files. For cultural heritage organizations, the management problem is no longer digitization alone. It is deciding which evidence can be trusted, which materials should be processed first, and which claims are sufficiently supported for preservation, access, exhibition, education, tourism, and resource allocation. Multimodal methods broaden the evidence base, but they become organizationally useful only when access, provenance, processing burden, and analytical completion are made explicit (Smits & Wevers, 2023; Tiribelli et al., 2024).

This problem is examined through Competitive Intelligence (CI). CI has developed from competitor monitoring and environmental scanning into an ethical, systematic, and decision-oriented process that connects intelligence needs, collection, analysis, communication, use, and feedback (Prescott, 1999; Calof & Wright, 2008; Dishman & Calof, 2008). Information becomes intelligence only when it is evaluated against a defined decision. Herring (1999) places the decision-maker's intelligence needs at the beginning of the process, while Pellissier and Nenzhelele (2013) emphasize the cyclical relationship among planning, collection, analysis, dissemination, organizational use, and feedback. This process perspective is distinct from text mining or descriptive reporting.

For cultural heritage institutions, relevant decisions include which collections require additional documentation, which digital assets justify further processing, where technical investment is needed, and which materials can responsibly support public interpretation. These institutions operate under budget constraints, changing technologies, regulatory duties, competition for public attention and funding, and expectations of social value. In this setting, sustainable competitive advantage refers to a durable organizational capability to convert heterogeneous heritage evidence into timely, defensible, and reusable knowledge while protecting authenticity and public value. The present study treats that capability as a theoretical organizational outcome, not as an empirically measured performance effect (Barney, 1991; Ranjan & Foropon, 2021).

Information asymmetry is central to this problem. Dataset creators may understand the file structure, while heritage managers, funders, researchers, and public users see only the repository description. A decision-maker may therefore assume that a multimodal deposit is immediately accessible or already suitable for analysis. Strategic factor market research shows that unequal access to relevant information can distort resource decisions (Makadok & Barney, 2001). In heritage analytics, provenance records, file inventories, evidence-status labels, and human oversight can reduce this asymmetry. Without these controls, the label “multimodal” may imply a level of analytical completion that the evidence does not support.

Chinese ancient architecture provides a suitable context because meaning is distributed across visual form, building type, ritual function, dynastic sequence, regional tradition, structure, and decoration. The Chinese ancient building multimodal dataset deposited in Harvard Dataverse contains image-text archives, videos, style-model files, and documentation (Li et al., 2024a, 2024b). It enables a CI-centered question that is organizational as well as technical: what decisions can be supported by the evidence that is actually accessible, and which decisions must be deferred until image-level and stakeholder validation are available?

The research problem is not whether the repository is multimodal, but whether repository evidence can be converted into reproducible management actions without being mistaken for a fully operating intelligence function. The study therefore asks: Which repository components are ready for immediate analysis, which should enter a controlled processing pilot, and which decisions must be withheld? Its objective is to test a pre-deployment CI protocol at repository level. The study contributes by placing the classical CI cycle at the center of the design, computing an empirical Repository Processing Readiness Index for the four modalities, applying role-based go/defer and stop rules, and specifying the evidence required for later image-level and institutional testing. The unit of inference remains the repository; no claim is made about asset-level cultural value or organizational performance.

2 THEORETICAL FRAMEWORK

2.1 Historical Evolution and Epistemological Foundations of Competitive Intelligence

The evolution of CI can be understood as a shift from ad hoc competitor data collection to an institutionalized intelligence function. Prescott (1999) describes its movement toward systems designed for managerial action, while Calof and Wright (2008) locate CI across practitioner, academic, and interdisciplinary traditions. Du Toit (2015) further shows that the field draws on information science, strategic management, marketing, economics, and organizational learning. CI is therefore neither a software product nor a synonym for data analysis. It is a knowledge-producing process whose quality depends on relevance, reliability, timeliness, ethical conduct, and usefulness to a decision.

Epistemologically, CI turns observations into evaluated meaning. Data describe objects or events; information organizes those data; intelligence interprets their implications for a defined organizational choice. Herring's (1999) key intelligence topics make this distinction operational by asking what decision must be made, what early warning is needed, and which actors or developments should be monitored. Planning therefore precedes collection. A repository should not be mined simply because it is available; its contents should be examined in relation to a stated intelligence requirement.

Environmental scanning broadens one's focus from just one competitor to the technological, regulatory, social, cultural, institutional and stakeholder environments that might influence an organization's options. Wright et al. (2002) demonstrate the varying levels of maturity, structure and managerial use of CI practices, and Dishman and Calof (2008) argue that collection and analysis must be linked to strategy formulation. Cultural heritage organisations are facing a range of environments that include conservation policy, changes in tourism demand, digital platforms, funding opportunities, copyright and constraints, scholarly standards, and changing audiences' attitudes.

The longer-term output of this process is strategic intelligence: an interpretation of external conditions and internal capabilities that helps decision-makers anticipate change and compare options. In the present study, file inventories, semantic indicators, and later multimodal features have value only when they answer an intelligence need, reduce uncertainty, and can be communicated to a responsible organizational user. This requirement links technical analysis to decision readiness rather than to data production alone.

2.2 The Competitive Intelligence Cycle as a Pre-Deployment Protocol

The CI cycle used in this study comprises nine connected functions. Intelligence planning defines the question, intended user, time horizon, and evidence threshold. Environmental scanning identifies technological, cultural, institutional, regulatory, and stakeholder signals. Intelligence collection gathers lawful, traceable, and decision-relevant evidence. Intelligence analysis evaluates quality, identifies patterns, and interprets implications. Intelligence dissemination presents results in a form suited to the intended audience. Intelligence utilization concerns the use of those results in prioritization, policy, budgeting, conservation, communication, education, tourism, or partnerships. Feedback intelligence captures user judgments and changing needs. Decision intelligence compares alternatives, constraints, and consequences. Intelligence governance operates across all stages through provenance, accountability, access control, ethics, validation, security, and human oversight (Herring, 1999; Pellissier & Nenzhelele, 2013).

The cycle is recursive rather than linear. Dissemination can expose new gaps, utilization can reveal that an indicator is not decision-relevant, and feedback can alter collection priorities. A mature intelligence capability therefore requires procedures, skilled analysts, suitable technology, organizational understanding, and an auditable record (Fleisher & Bensoussan, 2015; Ranjan & Foropon, 2021). In this article, however, the cycle serves as a pre-deployment protocol rather than as evidence that all nine functions operated within an institution. Three validation levels are distinguished: repository validation; multimodal record validation; and organizational validation. Only repository validation is examined empirically. Image-text analysis and organizational use remain separate empirical stages.

2.3 Multimodal Semantic Analytics in Cultural Heritage

As the number of texts being digitized grows, text-based approaches remain dominant in digital humanities, and more recently, digital humanities is shifting away from digitization towards computational interpretation (Luhmann & Burghardt, 2022; Joo et al., 2022; Suissa et al., 2022). The visual analysis is therefore not considered an auxiliary method, but a substantive interpretative method. The contribution of built and visual heritage to the learning of computer vision, GeoAI, building classification and large-scale studies on scene detection demonstrates the potential for learning from the built and visual heritage, and the need for domain labels and expert interpretation (Smits & Wevers, 2021, 2023; van Noord, 2022; Mishra & Lourenço, 2024; Pierdicca & Paolanti, 2022; Siountri & Anagnostopoulos, 2023; Trier et al., 2021; Wevers et al., 2022). Human-centred and ethical research also demands automated output to be accountable to provenance, accessibility, and meaning in culture (Pavlidis, 2022; Pisoni et al., 2021; Tiribelli et al., 2024).

Knowledge graphs enable linking of records, entities, periods, places and materials, but it is quality representation, representation completion techniques and representation interpretation by users that make the difference (Ji et al., 2022; Zamini et al., 2022; Khoo, 2024). Heritage applications illustrate various ways of using linked data and machine learning to enhance metadata and link together scattered records (Bobasheva et al. 2022; Debruyne et al. 2022; Fan et al. 2023; Huang et al. 2023; Tzitzikas et al., 2022). Record alignment and feature discovery tools are vision-language learning, large-scale image-text resources, segmentation, image generation, and cross-modal extraction (Jia et al., 2021; Radford et al., 2021; Schuhmann et al., 2022; Kirillov et al., 2023; Rombach et al., 2022; Yang et al., 2022; Yu et al., 2020). Those techniques, however, assume that records of images and texts are both available and consistent. These methods are not found in the executed run, and therefore are not part of the current results, but they are part of the later analytical stages.

2.4 Information Asymmetry, Intelligence Capability, and Sustainable Advantage

Information asymmetry arises when stakeholders possess unequal knowledge about whether heritage data exist, what they contain, how accessible they are, and what decisions they can support. A repository description may signal rich multimodality without disclosing extraction costs, record linkage, or annotation quality. CI capability reduces this asymmetry through repeated scanning, source assessment, provenance control, confidence statements, and explicit disclosure of what is known and what remains uncertain. This is an organizational capability because it depends on routines and governance, not on an algorithm alone.

Sustainable advantage is more likely to arise from a valuable and difficult-to-replicate capability than from possession of a public dataset itself (Barney, 1991). For heritage organizations, the relevant capability is the repeated conversion of cultural data into credible priorities: which collections need more documentation, which assets justify further digitization, which records can support public interpretation, and where conservation or communication resources should be directed. The present study does not test performance effects. It evaluates the decision readiness of the governed analytical process that would support such choices.

2.5 Research Gap and Analytical Propositions

The analysis is guided by three gaps. The theoretical gap is the limited integration of classical CI processes with multimodal Digital Humanities. The methodological gap is the frequent movement from repository descriptions to content-level claims without a transparent record of what was accessed and analyzed. The practical gap is the lack of reproducible rules that turn repository evidence into bounded management actions. Proposition 1 states that multimodal heritage data can enter a CI process only when image, text, metadata, and provenance are accessible and alignable. Proposition 2 states that repository processing readiness can be quantified from traceability, direct access, analytical completion, computational feasibility, and decision relevance. Proposition 3 states that role-based decision rules can produce defensible preparation actions, while organizational utilization and outcomes require observation in an institutional setting.

Figure 1 presents the Data-to-Intelligence logic used in the study. It links the dataset foundation to semantic and visual-structural layers, multimodal fusion, and CI outputs. The figure separates components supported by the executed evidence from components that require image-level or organizational validation. It should therefore be read as a staged decision-readiness framework rather than as a claim that the entire CI cycle has already been validated in practice.

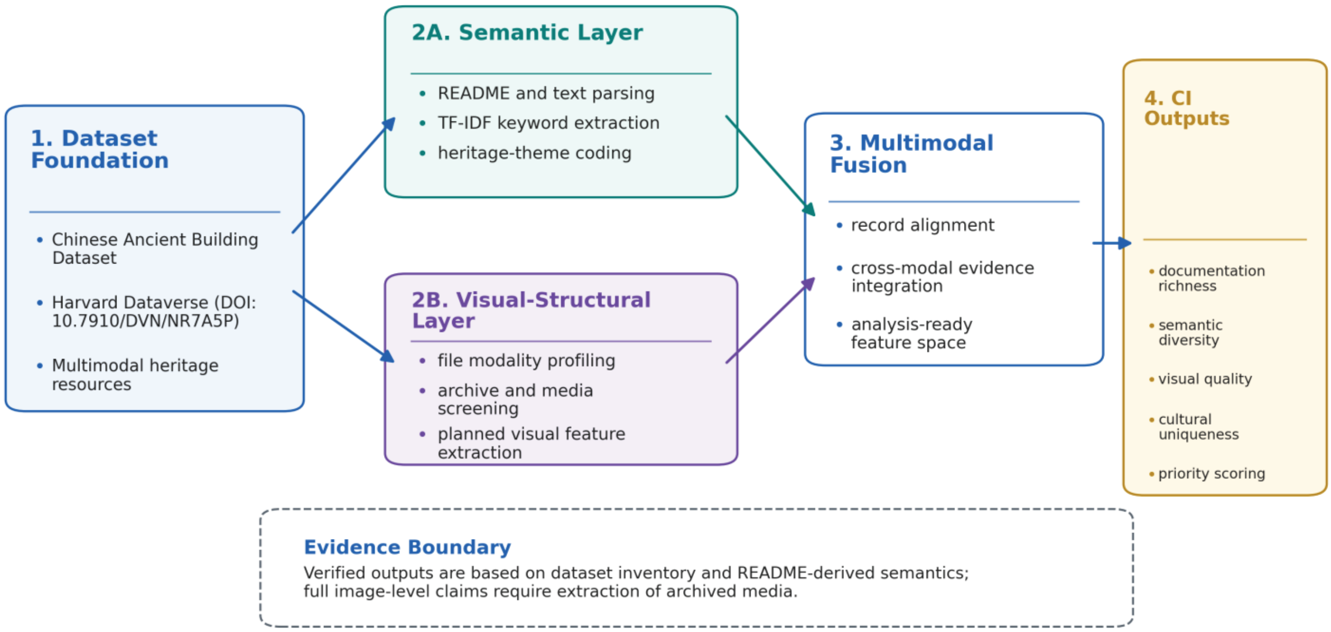

Figure 1. Pre-deployment Competitive Intelligence protocol for multimodal cultural heritage repositories.

3 METHOD

3.1 Research Design, Data Source, and Units of Analysis

This study uses a quantitative-descriptive, documentary, pre-deployment protocol design. The empirical purpose is to test repository processing readiness and the consistency of bounded decision rules; it does not assess organizational performance or complete multimodal interpretation. The primary units are the 35 deposited files and the accessible README. A second analytical unit is the modality group - text, archive, video, or model - used to calculate repository processing readiness. Individual heritage images and their annotations remain the required units for later asset-level analysis.

The source is the Chinese ancient building multimodal dataset (DOI 10.7910/DVN/NR7A5P), described by Li et al. (2024a, 2024b). The repository was selected because it combines Chinese cultural heritage content with image-text archives, videos, style-model files, and documentation. Table 1 distinguishes executed variables, modality-level readiness indicators, and measures reserved for record-level analysis. This separation prevents a repository score from being interpreted as a cultural-value score.

Table 1. Variable Definitions and Repository-Readiness Operationalization

Variable	Role	Definition	Measurement	Data source/ column	Intelligence purpose or status
file_id	Identifier	Unique Dataverse file identifier	Numeric ID	Dataverse metadata	Traceability and source control
filename	Descriptor	Deposited file name	Text string	Dataverse metadata	Collection mapping and retrieval
modality	Scanning variable	Video, model, archive, or text classification	Categorical	Filename, extension, MIME type	Environmental scanning and collection route
filesize_mb	Feasibility variable	Storage burden in megabytes	Continuous	Dataverse metadata	Infrastructure, cost, and sequencing decisions
download_ status	Collection variable	Whether a selected file was obtained	Categorical	Processing log	Evidence-access status
word_token_ count	Text indicator	Tokens in processed documentation	Count	README parsing	Documentation richness
heritage_ theme_ counts	Semantic indicators	Counts of defined architectural and cultural terms	Counts	Dictionary coding	Initial cultural-semantic signal
tfidf_score	Analytical indicator	Relative term weight in accessible text	Continuous	TF-IDF output	Documentation-level intelligence analysis
repository_ decision_ case	Decision-mapping variable	Bounded repository-management question, possible user, and proposed action linked to observed evidence	Structured mapping	Decision-readiness protocol	Shows decision relevance without claiming observed organizational use
evidence_ grade	Governance variable	Grade A observed repository evidence, Grade B reasoned recommendation, or Grade C validation-dependent application	Three-level classification	Method log	Separates repository, multimodal, and organizational evidence
CI priority index	Deferred outcome	Weighted completeness, documentation, visual quality, diversity, and uniqueness	Planned 0-100 score	Formula specification	Not computed; belongs to multimodal record validation

3.2 Inclusion, Exclusion, and Intelligence Collection Procedures

All files listed by the Dataverse API were included in the inventory and classified by filename, extension, MIME type, modality, and size. At the content-processing level, only directly accessible text was eligible for semantic analysis. Videos, model weights, and compressed archives were retained in the inventory but deferred from content analysis because they required additional storage, extraction time, or specialized processing. This rule was based on feasibility, not substantive irrelevance. It preserves the collection record while preventing unprocessed media from being presented as analyzed evidence.

3.3 Analytical Procedures

The executed procedures comprised: (a) Dataverse inventory retrieval; (b) file-type and storage profiling; (c) download logging; (d) README parsing; (e) Chinese-English tokenization; (f) TF-IDF term weighting; and (g) dictionary-based coding of roof/eaves, dougong and timber structure, temple or religious space, hall/gate/courtyard, dynastic period, regional heritage, and decorative style. Descriptive counts were appropriate because the empirical material consisted of 35 file records and one processed document. Inferential tests, regression, structural equation modeling, predictive machine learning, topic modeling, and correlation analysis were not used because the required record-level observations were unavailable.

At repository level, the Repository Processing Readiness Index (RPRI) was calculated as 20(T + A + C + F + D), where each component ranges from 0 to 1. T denotes metadata traceability, A direct content access, C analytical completion, F storage feasibility, and D decision relevance. Storage feasibility was calculated as 1 - ln(1 + mean file size) divided by the maximum modality value. Decision relevance was coded before scoring from the declared dataset role: archive = 1.00, video = 0.60, text = 0.50, and model = 0.20. A sensitivity specification assigned 40% to decision relevance and 15% to each remaining component. The index measures immediate processing readiness, not cultural importance, preservation urgency, or organizational performance.

The separate asset-level CI Priority Index shown in Figure 2 remains reserved for image-annotation records after extraction. It combines multimodal completeness, documentation richness, visual quality, semantic diversity, and cultural uniqueness. No asset received that score because the required record-level inputs were unavailable. Keeping the modality-level RPRI separate from the asset-level index prevents processing feasibility from being treated as heritage value.

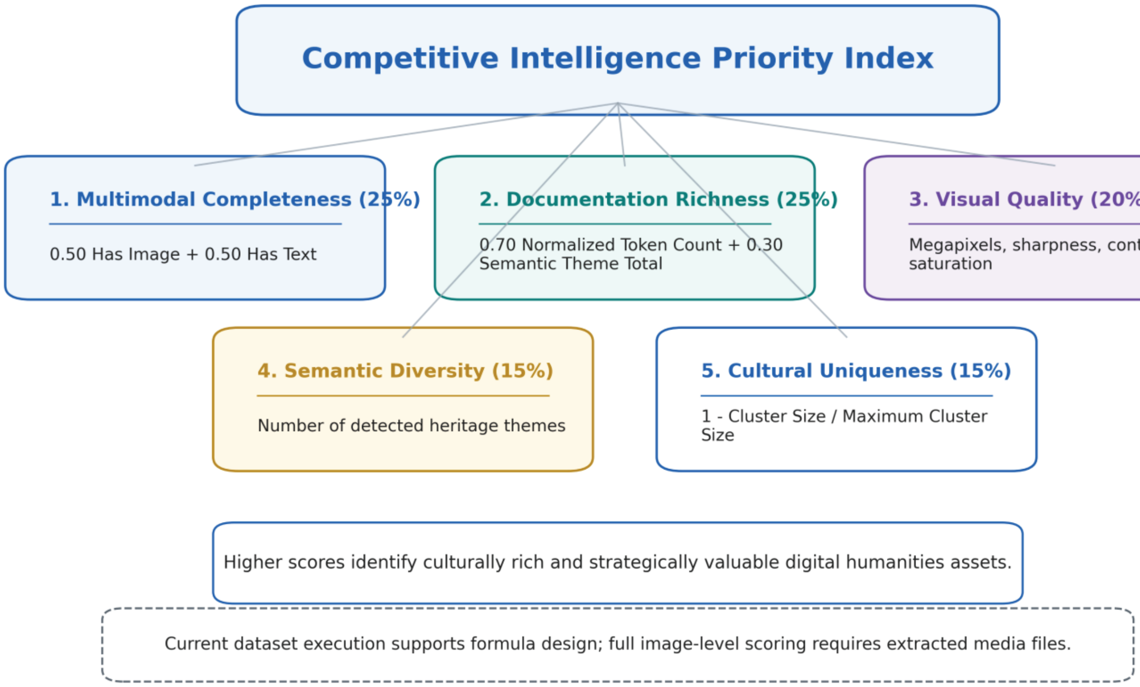

Figure 2. Asset-level CI Priority Index components reserved for record-level testing.

3.4 Processing Readiness, Structured Decision Simulation, Governance, and Reproducibility

The CI cycle was used to structure a pre-deployment protocol. Four bounded questions guided the test: whether any component is ready for immediate analysis; which file group should enter the first controlled pilot; which claims require a stop rule; and what evidence is needed before asset-level or institutional decisions. Grade A denotes directly observed repository evidence, Grade B a rule-based recommendation derived from Grade A evidence, and Grade C an application requiring image-level or organizational testing.

The observed scores were then applied in a structured decision simulation with four predefined roles. The heritage program manager received a conditional go for a bounded archive-extraction pilot rather than full deployment. The data steward received a requirement to preserve file identities, document extraction overhead, verify checksums, and test image-annotation pairing. The curator received a stop rule against asset ranking until annotations are validated. The communication or tourism lead received a restriction against presenting README themes as collection-wide cultural findings. This exercise demonstrates protocol-level utilization because the evidence produces explicit actions; it does not show that a real institution adopted those actions or experienced an outcome.

Intelligence governance was treated as a cross-cutting control. The rules were source traceability, lawful public-data use, preservation of original file identities, explicit separation of observed and planned variables, non-imputation of unavailable media content, human interpretation of cultural categories, evidence grading, role assignment, and stop rules for unsupported claims. These controls address provenance, bias, accountability, access, and the risk that technical outputs could be treated as authoritative heritage judgments (Pisoni et al., 2021; Tiribelli et al., 2024).

Validity was assessed at repository level through concept-indicator alignment and weight sensitivity. File size measures processing burden, not cultural value; README theme counts measure documentation signals, not the meaning of the archive. The equal-weight and decision-relevance-weighted specifications produced the same stable conclusions: text was the only immediately analyzable modality, model files ranked last for heritage analysis, and archives and videos required staged processing. Reliability was supported by rule-based classification, explicit formulas, a fixed coding dictionary, retained intermediate outputs, and a traceable role-action protocol. These controls establish reproducibility of the repository test, not validity of image-level priorities or organizational performance. Figure 3 summarizes the executed workflow and distinguishes completed evidence from deferred multimodal processing.

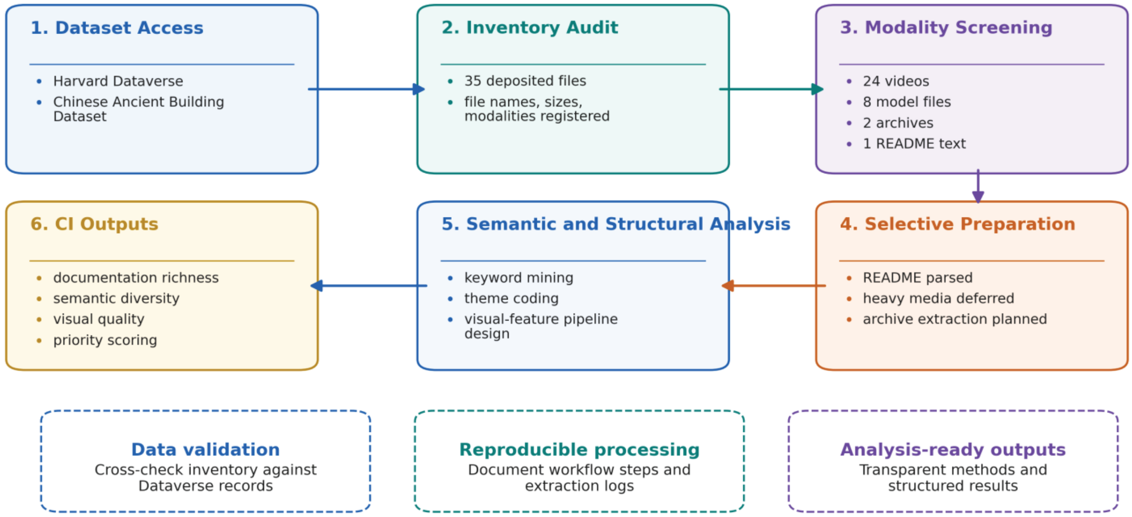

Figure 3. Repository protocol, readiness scoring, and empirical evidence boundary.

4 RESULTS AND DISCUSSION

4.1 Environmental Scanning and Collection Readiness

The repository inventory identified 35 deposited files totaling 18,816.252 MB. Twenty-four videos represented 68.6% of files but only 10.7% of storage. Eight style-model files represented 22.9% of files and 69.4% of storage. Two compressed archives represented 5.7% of files and 19.9% of storage. The single README file was negligible in size but was the only content directly processed. The deposit is therefore multimodal at source level, while the executed evidence remains at repository and documentation level.

Table 2 reports this distribution. In CI terms, file count and storage burden are environmental-scanning inputs because they can inform collection strategy, infrastructure, cost, sequencing, and timing. On the present evidence, the defensible recommendation is to inspect the two archives in a controlled pilot before authorizing asset-level interpretation. Whether an institution accepts or acts on that recommendation was not tested.

Table 2. Dataset Inventory by Modality and Intelligence-Collection Implication

Modality	N files	Total size (MB)	Share of files	Competitive Intelligence interpretation
Video	24	2013.914	68.6%	Repository coverage signal; content analysis requires frame extraction or scene coding
Model	8	13051.374	22.9%	Largest storage burden; technical resource, not a heritage observation
Archive	2	3750.958	5.7%	Potential image-text core; extraction and pairing required
Text	1	0.006	2.9%	Directly accessible documentation used for the executed semantic audit
Total dataset inventory	35	18816.252	100.0%	Source-level multimodality established; analyzed multimodality remains limited

4.2 Semantic Analysis and Information Asymmetry

The README audit produced 5,872 characters and 561 tokens, with no Chinese characters. This finding does not contradict the Chinese context of the dataset; it shows that the directly accessible documentation is English and cannot substitute for the archived annotations. Table 3 presents the observed semantic signals. The audit reduces information asymmetry by showing what the accessible document contains, while also exposing the larger gap between the repository description and the record-level cultural evidence held inside the archives.

Table 3. README Semantic Statistics and Strategic Interpretation

Indicator	Observed value	Strategic interpretation
README files processed	1	Directly accessible documentation; not the full annotation corpus
Character count	5872	Limited documentation volume
Word/token count	561	Small text corpus suitable for descriptive semantic analysis
Chinese character count	0	Accessible documentation is English-language
Roof/eaves theme count	3	Architectural vocabulary signal
Dougong/timber theme count	1	Structural vocabulary signal
Dynastic period theme count	5	Historical-period documentation signal
Regional heritage theme count	1	Limited regional/contextual vocabulary
Decoration/artistic style count	5	Strongest heritage-theme signal with dynastic terms
Top TF-IDF terms	safetensors, concept, can, which, sketch, Chinese, design, buildings, front, urban, drawings, drawing	Technical and style-oriented documentation dominates the accessible text

Figure 4 ranks the dominant TF-IDF terms. “Safetensors,” “concept,” “sketch,” and “design” receive greater weight than most heritage-domain terms. The README therefore functions mainly as technical and style-model documentation. For intelligence analysis, it clarifies the production environment of the deposit, but it is insufficient for judging the historical importance or conservation priority of individual buildings.

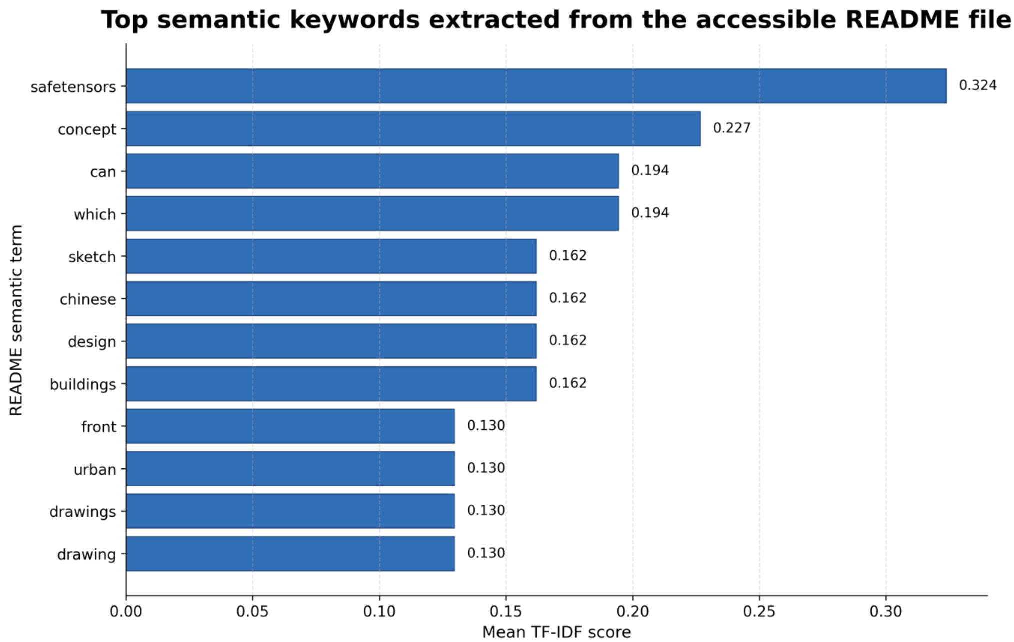

Figure 4. TF-IDF keyword profile of the accessible README file.

Figure 5 shows the dictionary-based heritage signals. Dynastic-period and decorative-style references were most frequent, followed by roof/eaves terminology; temple/religious and hall/gate/courtyard terms were absent. These zeros indicate absence from the processed README, not absence from the underlying collection. Treating that distinction as a formal stop rule is an intelligence-governance requirement because documentation gaps must not be converted into claims about cultural reality.

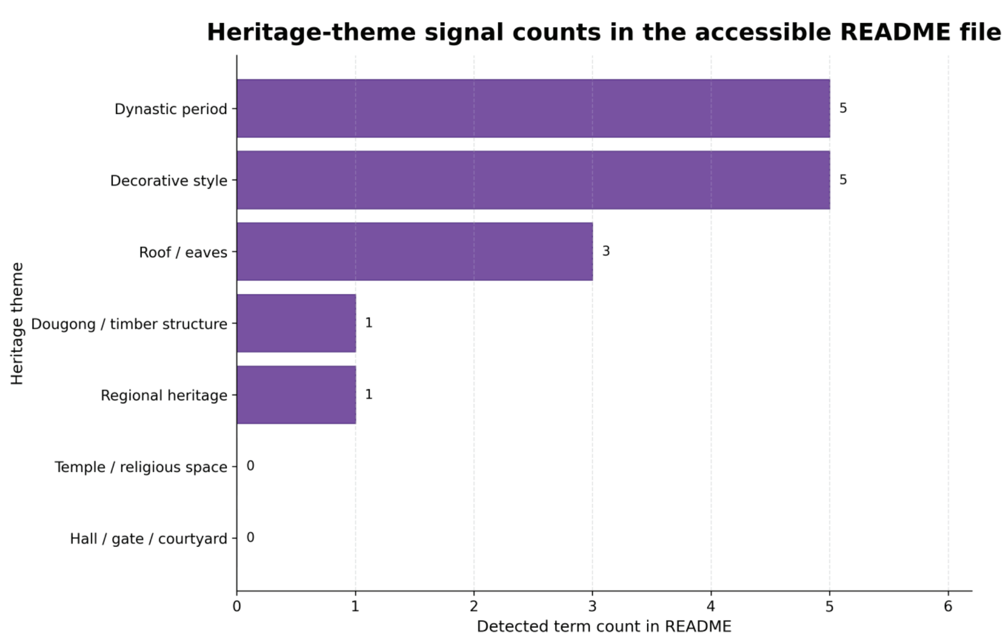

Figure 5. Heritage-theme signals detected in the accessible README file.

The findings confirm that multimodal heritage intelligence requires domain context and source evaluation. Knowledge graphs and vision-language methods can enrich cultural records, but they cannot compensate for unextracted media or narrow documentation (Fan et al., 2023; Huang et al., 2023; Smits & Wevers, 2023). At this stage, the semantic audit is a repository-level early-warning signal: the accessible text is technically informative but too narrow for asset-level interpretation. It does not demonstrate organizational intelligence use.

4.3 Empirical Processing Readiness and Structured Decision Outputs

Table 4. Repository Processing Readiness Scores and Protocol Actions

Modality and observed basis	RPRI result	Protocol action and intended role
Text: 1 README; 0.006 MB; directly accessible and analyzed	90.0 equal-weight; 80.0 sensitivity	Retain as documentation evidence only; prohibit asset ranking. Curator and metadata manager.
Archive: 2 files; 3,750.958 MB; declared image-text core; not extracted	40.0 equal-weight; 55.0 sensitivity	Conditional go for a bounded extraction and pairing pilot. Heritage program manager, data steward, and curator.
Video: 24 files; 2,013.914 MB; metadata available; frames not analyzed	40.2 equal-weight; 45.2 sensitivity	Plan frame extraction after archive structure and pairing are verified. Data steward and technical analyst.
Model: 8 files; 13,051.374 MB; technical weights rather than heritage observations	24.4 equal-weight; 23.3 sensitivity	Defer from direct heritage interpretation; inspect model documentation only. Technical lead and governance reviewer.

The RPRI converts observed repository conditions into a modality-level empirical output. Text scored 90.0 because it was traceable, directly accessible, analyzed, and computationally light. Video scored 40.2, archives 40.0, and model files 24.4. These values indicate processing readiness rather than heritage significance. Under the sensitivity specification, text scored 80.0, archives 55.0, video 45.2, and model files 23.3. The stable findings are that the README is the only immediately analyzable component and that model weights should not lead the heritage-analysis sequence. Giving greater weight to decision relevance places the image-text archives ahead of video, supporting an archive-first pilot.

Table 5. Competitive Intelligence Functions, Demonstrated Protocol Outputs, and Remaining Evidence

CI function	What the current study provides	Evidence status	Evidence required beyond this study
Intelligence planning	Four bounded questions, intended roles, evidence grades, and decision thresholds	Demonstrated within the repository protocol	Confirm institutional priorities with heritage managers
Environmental scanning	Repository modalities, formats, storage burden, access conditions, and documentation scope	Demonstrated for the selected repository	Compare signals across repositories and institutional settings
Intelligence collection	API metadata and accessible README collected; deferred file identities preserved	Partially demonstrated at repository level	Extract and verify image-annotation pairs

CI function	What the current study provides	Evidence status	Evidence required beyond this study
Intelligence analysis	Inventory statistics, TF-IDF, theme counts, RPRI scores, and sensitivity analysis	Demonstrated at repository and documentation level	Compute image-level features, clusters, and asset-level CI priorities
Intelligence dissemination	Evidence-status tables, readiness scores, role-specific brief, confidence grades, and stop rules	Demonstrated as a structured research brief; external comprehension not tested	Deliver the brief to intended institutional users and test comprehension
Intelligence utilization	Conditional go/defer and stop rules applied in a four-role decision simulation	Demonstrated at protocol level; institutional adoption not observed	Conduct an institutional case and document acceptance, rejection, or action
Intelligence governance	Provenance, evidence grades, role accountability, claim controls, human oversight, and audit trail	Demonstrated as protocol controls; institutional governance not tested	Audit governance performance during institutional use
Feedback intelligence	Alternative weighting tested to assess stability of the processing sequence	Internal sensitivity feedback demonstrated; external user feedback not collected	Obtain and analyze feedback from intended users
Decision intelligence	RPRI-based comparison generated a bounded archive-first pilot and explicit defer rules	Demonstrated as a protocol decision output; organizational consequences not observed	Track resource commitments, implementation, and consequences in a live case

Table 5 separates functions demonstrated within the protocol from functions requiring a live institutional case. Planning, scanning, collection, analysis, governance, and modality-level scoring were performed by the research team. Dissemination took the form of a decision brief and role-specific outputs. Utilization was demonstrated only inside the structured simulation through conditional go/defer and stop rules. Weight sensitivity provided internal feedback on the stability of the processing sequence. External user feedback, institutional adoption, resource commitments, and decision consequences were not observed.

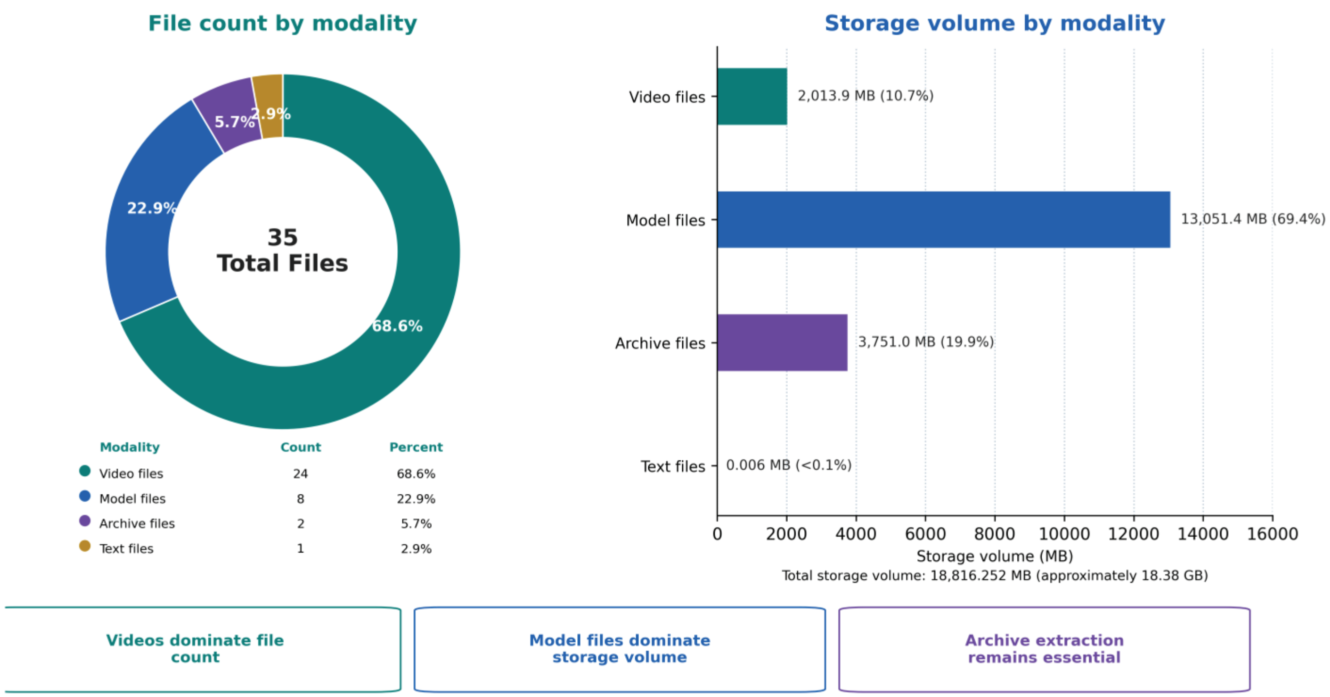

Figure 6. Modality composition and storage profile of the Chinese ancient building dataset.

4.4 Empirical Boundary and Required Validation Stages

The blank image-feature, multimodal-pair, clustering, cluster-summary, asset-priority, and location-intelligence sheets remain an audit trail of record-level stages that were not executed. They are not the endpoint of the present analysis. The repository-level endpoint is the computed RPRI, its sensitivity test, and the role-based decision outputs. Image-level cultural interpretation still requires extraction and validation of the archived records.

Figure 7 summarizes the three-stage evidence path. The present study completes repository validation through inventory, documentation audit, modality-level readiness scoring, sensitivity analysis, and structured decision simulation. Multimodal record validation requires image-annotation pairing, feature analysis, clustering, and asset-level scoring. Organizational validation requires delivery to intended users, observed adoption or rejection, feedback, resource commitments, and recorded consequences. The latter two stages remain distinct from the completed repository protocol test.

The repository evidence supports four explicit preparation decisions: retain the README as documentation evidence only; authorize a bounded pilot on the two archives; plan video frame extraction after archive structure and pairing are verified; and defer model weights from direct heritage interpretation. It cannot support visual-quality rankings, conservation priorities, tourism recommendations, or claims of organizational effectiveness.

Organizational validation remains a separate research stage. It begins when intended users receive the outputs and their responses are recorded. A later institutional case should measure comprehension, changes to priorities or weights, acceptance or rejection of recommendations, action taken, resource commitment, and consequences for decision timeliness, consistency, and traceability. The present study supplies the protocol and executable decision rules for that test rather than claiming those outcomes.

	Current accessibility	Current analysis use	Immediate contribution	Next-step requirement
README text	Directly accessible	Keyword extraction	Initial semantic evidence	Expanded annotation parsing
Videos	Metadata accessible	Inventory only	Coverage of heritage sites	Frame extraction / OCR / scene coding
Model files	Metadata accessible	Infrastructure evidence	Shows visual-style resources	Model-card interpretation
Archives	Deposited but heavy	Deferred extraction	Potential multimodal core	Unzip and process image-text assets

Interpretive implication
The dataset strongly supports a validated analysis-ready framework, but full multimodal claims depend on future extraction of archived media and record-level image-text linkage.

Figure 7. Completed repository protocol and the two remaining validation stages.

5 FINAL CONSIDERATIONS

This pre-deployment study examined whether a multimodal Chinese cultural heritage repository can produce defensible CI preparation decisions. The work verified 35 files, quantified modality and storage burdens, processed one accessible README, calculated modality-level readiness scores, tested an alternative weighting scheme, and executed role-based go/defer and stop rules. It did not conduct image-level interpretation or observe institutional outcomes.

The theoretical contribution is to position CI as the organizing logic for a staged evidence process rather than as a label applied after Digital Humanities analysis. The framework connects intelligence needs, information asymmetry, governance, processing readiness, dissemination, protocol-level utilization, feedback requirements, and decision thresholds. It also distinguishes a demonstrated repository protocol from later organizational capability.

The methodological contribution is an evidence-controlled repository protocol with a computed RPRI, sensitivity analysis, evidence grades, role assignments, and stop rules. The study distinguishes deposited, accessible, analyzed, and decision-ready evidence. This provides a reproducible basis for selecting the first processing pilot without presenting storage readiness as cultural value or simulated action as institutional performance.

For practice, the paper produces four direct outputs: use the README only for documentation-level intelligence; begin with a controlled extraction of the two image-text archives; schedule video processing after archive pairing is verified; and keep model weights outside asset-level heritage interpretation unless their documentation establishes a specific analytical purpose. These outputs are generated by the protocol and can be taken into an institutional validation case.

The decision boundary is precise. The completed analysis supports repository preparation, sequencing, documentation control, and governance. It does not support building-level conservation rankings, regional comparisons, tourism recommendations, organizational effectiveness, or sustainable competitive advantage. Those claims require validated image-annotation records and observed decisions by responsible institutional users.

The article should therefore be read as a pre-deployment protocol study with an empirical repository-level index and structured decision test. Its contribution is the reproducible conversion of repository evidence into bounded actions and explicit thresholds for the next two validation stages. It does not claim that the complete CI cycle has operated inside a cultural heritage organization.

5.1 Limitations and Future Research

The principal limitation is the absence of extracted image-text records and participating organizations. Consequently, the study contains no cross-modal embeddings, clusters, asset-level cultural scores, stakeholder adoption records, or decision outcomes. The accessible README is English-language technical documentation and may not represent the archived Chinese annotations. The RPRI is a modality-level processing measure based on transparent rule coding; it must not be interpreted as cultural value.

Future work must cross two empirical thresholds. First, the archives should be extracted, image-annotation pairs validated, and visual features, clusters, and the asset-level CI Priority Index calculated with expert checks for cultural bias. Second, the resulting brief should be tested with heritage managers, curators, data stewards, and education or tourism leads. Comprehension, changes to priorities or weights, acceptance or rejection, actions taken, resource commitments, and decision consequences should be recorded.

Only after those stages should cross-institutional research assess repeatability across collections and governance settings. Claims of organizational effectiveness or sustainable advantage require that evidence.

DECLARATIONS

Funding information: was not provided.

Conflict-of-interest information: was not provided.

Ethics approval information: was not provided. The analysis used a public dataset inventory and non-human documentation files; no direct participant data were analyzed.

Data availability: The source dataset is available through Harvard Dataverse at https://doi.org/10.7910/DVN/NR7A5P. The reported results comprise the Dataverse inventory, the accessible README semantic audit, modality-level processing-readiness scores, a sensitivity specification, and structured decision-rule outputs derived from those records.

REFERENCES

Barney, J. B. (1991). Firm resources and sustained competitive advantage. *Journal of Management, 17*(1), 99-120. https://doi.org/10.1177/014920639101700108

Bobasheva, A., Gandon, F., & Precioso, F. (2022). Learning and reasoning for cultural metadata quality: Coupling symbolic AI and machine learning over a semantic web knowledge graph to support museum curators in improving the quality of cultural metadata and information retrieval. *Journal on Computing and Cultural Heritage, 15*(3), 1-23. https://doi.org/10.1145/3485844

Calof, J. L., & Wright, S. (2008). Competitive intelligence: A practitioner, academic and interdisciplinary perspective. *European Journal of Marketing, 42*(7/8), 717-730. https://doi.org/10.1108/03090560810877114

Debruyne, C., Munnelly, G., Kilgallon, L., O'Sullivan, D., & Crooks, P. (2022). Creating a knowledge graph for Ireland's lost history: Knowledge engineering and curation in the Beyond 2022 project. *Journal on Computing and Cultural Heritage, 15*(2), 1-25. https://doi.org/10.1145/3474829

Dishman, P. L., & Calof, J. L. (2008). Competitive intelligence: A multiphasic precedent to marketing strategy. *European Journal of Marketing, 42*(7/8), 766-785. https://doi.org/10.1108/03090560810877141

Du Toit, A. S. A. (2015). Competitive intelligence research: An investigation of trends in the literature. *Journal of Intelligence Studies in Business, 5*(2), 14-21. https://doi.org/10.37380/jisib.v5i2.127

Fan, T., Wang, H., & Hodel, T. (2023). CICHMKG: A large-scale and comprehensive Chinese intangible cultural heritage multimodal knowledge graph. *Heritage Science, 11*(1), 115. https://doi.org/10.1186/s40494-023-00927-2

Fleisher, C. S., & Bensoussan, B. E. (2015). *Business and competitive analysis: Effective application of new and classic methods* (2nd ed.). FT Press.

Herring, J. P. (1999). Key intelligence topics: A process to identify and define intelligence needs. *Competitive Intelligence Review, 10*(2), 4-14. https://doi.org/10.1002/(SICI)1520-6386(199932)10:2%3C4::AID-CIR3%3E3.0.CO;2-C

Huang, Y., Yang, C., Gu, J., Jiang, P., Li, B., & Liu, Z. (2023). Using knowledge graphs and deep learning algorithms to enhance digital cultural heritage management. *Heritage Science, 11*(1), 204. https://doi.org/10.1186/s40494-023-01042-y

Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems, 33*(2), 494-514. https://doi.org/10.1109/TNNLS.2021.3070843

Jia, C., Yang, Y., Xia, Y., Chen, Y. T., Parekh, Z., Pham, H., Le, Q. V., Sung, Y. H., Li, Z., & Duerig, T. (2021). Scaling up visual and vision-language representation learning with noisy text supervision. *Proceedings of Machine Learning Research, 139*, 4904-4916. https://proceedings.mlr.press/v139/jia21b.html

Joo, S., Hootman, J., & Katsurai, M. (2022). Exploring the digital humanities research agenda: A text mining approach. *Journal of Documentation, 78*(4), 853-870. https://doi.org/10.1108/JD-03-2021-0066

Khoo, M. (2024). Knowledge graph visualization interface for digital heritage collections: Design issues and recommendations. *Information Technology and Libraries, 43*(1). https://doi.org/10.5860/ital.v43i1.16719

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W. Y., Dollár, P., & Girshick, R. (2023). Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 4015-4026). https://doi.org/10.1109/ICCV51070.2023.00371

Li, B., Feng, J., Yan, Y., Kou, G., Li, H., Du, Y., Wang, X., & Guo, K. (2024a). Building a Chinese ancient architecture multimodal dataset combining image, annotation and style-model. *Scientific Data, 11*(1), 1137. https://doi.org/10.1038/s41597-024-03946-1

Li, B., Feng, J., Yan, Y., Shi, Y., & Kou, G. (2024b). *Chinese ancient building multimodal dataset* [Data set]. Harvard Dataverse. https://doi.org/10.7910/DVN/NR7A5P

Luhmann, J., & Burghardt, M. (2022). Digital humanities: A discipline in its own right? An analysis of the role and position of digital humanities in the academic landscape. *Journal of the Association for Information Science and Technology, 73*(2), 148-171. https://doi.org/10.1002/asi.24533

Makadok, R., & Barney, J. B. (2001). Strategic factor market intelligence: An application of information economics to strategy formulation and competitor intelligence. *Management Science, 47*(12), 1621-1638. https://doi.org/10.1287/mnsc.47.12.1621.10245

Mishra, M., & Lourenço, P. B. (2024). Artificial intelligence-assisted visual inspection for cultural heritage: State-of-the-art review. *Journal of Cultural Heritage, 66*, 536-550. https://doi.org/10.1016/j.culher.2024.01.005

Pavlidis, G. (2022). AI trends in digital humanities research. *Trends in Computer Science and Information Technology, 7*(1), 13-18. https://doi.org/10.17352/tcsit.000048

Pellissier, R., & Nenzhelele, T. E. (2013). Towards a universal competitive intelligence process model. *South African Journal of Information Management, 15*(2), Article 567. https://doi.org/10.4102/sajim.v15i2.567

Pierdicca, R., & Paolanti, M. (2022). GeoAI: A review of artificial intelligence approaches for the interpretation of complex geomatics data. *Geoscientific Instrumentation, Methods and Data Systems, 11*(1), 195-218. https://doi.org/10.5194/gi-11-195-2022

Pisoni, G., Díaz-Rodríguez, N., Gijlers, H., & Tonolli, L. (2021). Human-centered artificial intelligence for designing accessible cultural heritage. *Applied Sciences, 11*(2), 870. https://doi.org/10.3390/app11020870

Prescott, J. E. (1999). The evolution of competitive intelligence: Designing a process for action. *Proposal Management*, Spring, 37-52.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. *Proceedings of Machine Learning Research, 139*, 8748-8763. https://proceedings.mlr.press/v139/radford21a.html

Ranjan, J., & Foropon, C. (2021). Big data analytics in building the competitive intelligence of organizations. *International Journal of Information Management*, *56*, 102231. https://doi.org/10.1016/j.ijinfomgt.2020.102231

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10684-10695). https://doi.org/10.1109/CVPR52688.2022.01042

Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., & Jitsev, J. (2022). LAION-5B: An open large-scale dataset for training next-generation image-text models. *Advances in Neural Information Processing Systems, 35*, 25278-25294.https://proceedings.neurips.cc/paper_files/paper/2022/hash/a1859debfb3b59d094f3504d5ebb6c25-Abstract-Datasets_and_Benchmarks.html

Siountri, K., & Anagnostopoulos, C. N. (2023). The classification of cultural heritage buildings in Athens using deep learning techniques. *Heritage*, *6*(4), 3673-3705. https://doi.org/10.3390/heritage6040195

Smits, T., & Wevers, M. (2021). The agency of computer vision models as optical instruments. *Visual Communication*, *21*(2), 329-349. https://doi.org/10.1177/1470357221992097

Smits, T., & Wevers, M. (2023). A multimodal turn in digital humanities: Using contrastive machine learning models to explore, enrich, and analyze digital visual historical collections. *Digital Scholarship in the Humanities*, *38*(3), 1267-1280. https://doi.org/10.1093/llc/fqad008

Suissa, O., Elmalech, A., & Zhitomirsky-Geffet, M. (2022). Text analysis using deep neural networks in digital humanities and information science. *Journal of the Association for Information Science and Technology*, *73*(2), 268-287. https://doi.org/10.1002/asi.24544

Tiribelli, S., Pansoni, S., Frontoni, E., & Giovanola, B. (2024). Ethics of artificial intelligence for cultural heritage: Opportunities and challenges. *IEEE Transactions on Technology and Society*, *5*(3), 293-305. https://doi.org/10.1109/TTS.2024.3432407

Trier, Ø. D., Reksten, J. H., & Løseth, K. (2021). Automated mapping of cultural heritage in Norway from airborne lidar data using Faster R-CNN. *International Journal of Applied Earth Observation and Geoinformation*, *95*, 102241. https://doi.org/10.1016/j.jag.2020.102241

Tzitzikas, Y., Mountantonakis, M., Fafalios, P., & Marketakis, Y. (2022). CIDOC-CRM and machine learning: A survey and future research directions. *Heritage*, *5*(3), 1612-1636. https://doi.org/10.3390/heritage5030084

van Noord, N. (2022). A survey of computational methods for iconic image analysis. *Digital Scholarship in the Humanities*, *37*(4), 1316-1338. https://doi.org/10.1093/llc/fqac003

Wevers, M., Vriend, N., & de Bruin, B. (2022). What to do with 2,000,000 historical press photos? The challenges and opportunities of applying a scene detection algorithm to a digitised press photo collection. *TMG Journal for Media History*, *25*(1), 1-25. https://doi.org/10.18146/tmg.815

Wright, S., Pickton, D. W., & Callow, J. (2002). Competitive intelligence in UK firms: A typology. *Marketing Intelligence & Planning*, *20*(6), 349-360. https://doi.org/10.1108/02634500210445400

Yang, L., Na, J. C., & Yu, J. (2022). Cross-modal multitask transformer for end-to-end multimodal aspect-based sentiment analysis. *Information Processing & Management*, *59*(5), 103038. https://doi.org/10.1016/j.ipm.2022.103038

Yu, J., Jiang, J., Yang, L., & Xia, R. (2020). Improving multimodal named entity recognition via entity span detection with unified multimodal transformer. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 3342-3352). https://doi.org/10.18653/v1/2020.acl-main.306

Zamini, M., Reza, H., & Rabiei, M. (2022). A review of knowledge graph completion. *Information*, *13*(8), 396. https://doi.org/10.3390/info13080396