



✦ Generated summary

bev art

White Paper

Navigating AI Use in Cultural Heritage

Balancing Innovation and Integrity

bev art

Artificial intelligence is entering the cultural heritage sector with growing momentum, offering conservators and collections professionals new capabilities in environmental monitoring, condition assessment, metadata management and public access.

But we must ask, in a field grounded by long-term stewardship and authenticity, how do we ensure accuracy, decide what's appropriate and adopt without comprising mission or trust?

The Juxtaposition of AI in Conservation Work

Cultural Heritage conservation is built on methodological rigor, material honesty and reversibility. In practice, this means decisions are documented, interventions are justified and the evidentiary chain from object to treatment to record is treated as part of the object's history. Conversely, AI has a different set of defaults: probabilistic reasoning, pattern recognition at scale, and outputs that are often accurate in aggregate but opaque in derivation.

That tension between conservation's demand for traceable reasoning and AI's tendency toward black-box conclusions is the central professional challenge of this moment. It is not a

matter of whether or not AI belongs in the cultural heritage sector because it is already here. Instead, we must ask how to engage with it deliberately and critically.

Real World Example: Trust, But Verify

A conservator is preparing a presentation and uses an LLM to help summarize and catalog an object. The model confuses it with a similarly named piece and gets the origin wrong. The conservator is busy, trusts the output, and the incorrect information makes it into the final presentation. No one questions it because they defer to the expert. The presentation is published and a member of the public catches the error and calls it out on social media.

What happened? **The important step of human verification was skipped — requesting sources, clicking through to read them and independently searching for verification outside the LLM**

Conservators occupy a distinct position in this conversation. Unlike administrators weighing operational efficiency, for example, they are directly accountable for the physical and chemical condition of objects. Errors in that domain are not easily corrected.

This raises the professional bar for AI adoption in preservation contexts considerably. A metadata error generated by an AI cataloging tool is recoverable. A treatment recommendation derived from a misread condition assessment is not. Conservators evaluating AI applications must apply an appropriate framework prior to adoption: is this tool's output reliable, traceable and within the bounds of accepted practice and what is the commitment to human review?

The Status of AI Adoption in Cultural Heritage

AI is being used in cultural heritage spaces in a number of ways – from administrative tasks like daily business communications to comprehensive data analysis.

One of the most mature and defensible applications of AI with cultural heritage conservation is environmental monitoring. Systems that integrate sensor data — temperature, relative humidity, light levels — and apply

predictive modeling to anticipate deterioration conditions offer conservators actionable information earlier than traditional monitoring approaches. When these systems are transparent in their methodology and integrated into existing preventive conservation workflows, they extend the conservator's capacity without replacing the conservator's judgment.

"When we built AI into our software, we knew our customers were already exporting data to external tools to save time on reports," said **Håvard Lystrup, Co-Founder & CEO of Bev/Art**. "We wanted to make that safer and smarter. The AI flags anomalies and surfaces risk but our customers' expertise determines what happens next. Technology should extend trust, not replace judgment."

AI-assisted condition analysis is a more developing area, with meaningful potential for large or distributed collections where manual survey is resource-intensive. The critical constraint here is validation: any AI-generated condition finding should be treated as a flag for conservator review, not a conclusion in itself.

Another area is collections

management. Object recognition, automated metadata generation and inventory support are reducing the administrative burden on conservation and collections staff in institutions that have piloted these tools carefully. The practical value is clearest in institutions managing significant cataloging backlogs, where AI can generate a working scaffold for human review rather than a finished record. The distinction matters. AI-generated metadata that enters collection management systems without conservator or registrar review introduces errors that compound over time.

Finally, AI is also being applied in automated transcription, multilingual translation and personalized visitor experience tools to expand public access to collections in ways that conservators are generally less directly involved in evaluating. Where conservators do have a stake is in how objects are represented in those contexts and whether interpretive AI applications maintain the distinction between documented fact and generated inference.

Navigating the Tensions of AI Use

Before adopting any framework for AI use, it's important to note the specific

challenges and considerations for conservators in cultural heritage.

- *Data Sensitivity & Ownership:* Integrating AI is complicated by culturally sensitive materials (e.g., sacred objects, human remains, repatriation-pending items) that existing data governance often fails to cover. Institutions must establish clear data governance policies regarding ownership and consent before high-resolution imaging or treatment records enter any AI pipeline.
- *Bias & Authenticity:* AI systems inherit biases from historical collection data, which can automate or amplify outdated frameworks without critical review. Conservators must scrutinize AI-assisted reconstructions and imaging applications, clearly maintaining the distinction between a documented condition and a hypothetical result.
- *Operational & Infrastructure Constraints:* Most cultural institutions are not resourced for the demands of AI adoption, including costs for staff capacity, training, legacy database integration and ongoing model evaluation. Conservators should be direct about these constraints in procurement and skeptical of solutions that minimize implementation cost.

- *Explainability*: Conservation standards require that treatment decisions be traceable and justifiable. AI tools that provide opaque outputs without interpretable reasoning are poorly suited for the field, as they fail to meet the evidentiary standard required for documentation. The onus is on conservators to meet these standards independent of AI results.

These challenges do not represent complete roadblocks but they do require navigating tradeoffs. Most of the tensions do not cleanly resolve themselves and require deliberate institutional understanding and positioning.

Central to all of them is automation versus human expertise. AI tools offer efficiency gains that are genuinely

attractive in under-resourced institutions, but the risk is that efficiency pressure gradually displaces human judgment from workflows where it remains essential. Conservators should be specific about where AI can appropriately reduce workload and where it cannot, and should document those boundaries explicitly in institutional AI policies.

A related tension point is transparency versus proprietary systems. Many AI tools operate on closed architectures that resist independent audit. Institutions that adopt mass market models, like ChatGPT or Claude, take on accountability for outputs they cannot fully evaluate. Where possible, preference should be given to systems whose methodologies are documented and whose outputs can be reviewed against known standards, such as custom institutional models.

Real World Example: The Importance of Disclosure

A conservator uses AI to help draft a grant narrative and shares it without disclosing that. A colleague reads it and something feels off — the voice isn't quite right, the phrasing a little too polished, the language familiar in a way they can't quite place. They don't say anything, but they notice. Elsewhere, the grantor has seen nearly identical phrasing in another institution's submission. No one makes an accusation. But something shifts in both relationships.

Are AI Bots Knocking Cultural Heritage Offline?

A report on the impact that bots building datasets for AI model training are having on online cultural collections in early 2025, by **Michael Weinberg**

[Read the full report](#)



Finally, scale versus specificity is a structural limitation. AI performs best on large, standardized datasets. Conservation collections are often small, materially heterogeneous and contextually specific in ways that general-purpose AI models are not designed to accommodate. As noted above, performance benchmarks from large-scale applications may not translate to specialized collection contexts, and institutions should require evidence specific to their use case before committing to adoption.

A Framework for Ethical AI Adoption

Several principles can guide conservators and institutions through AI adoption decisions.

- *Build cross-functional evaluation teams:* AI procurement decisions that exclude conservators, registrars and collections staff from the evaluation process are more likely to produce tools that fail in practice. The professionals closest to the collection are essential evaluators of whether a tool performs as claimed on real collection data.
- *Start with pilots.* Proof-of-concept implementations on bounded, lower-risk applications generate evidence that generalized vendor claims cannot provide. They also build institutional familiarity with AI systems in a context where failure is recoverable.
- *Require explainability as a baseline.* Any AI tool applied in a conservation context should be able to produce reasoning that a conservator can evaluate. If a vendor cannot demonstrate how their system reaches its conclusions, that is a disqualifying

limitation for preservation applications.

- *Establish data governance before deployment.* Institutions should understand what data AI tools consume, how it is stored, who has access to it, and under what terms it may be used for model training or other purposes. These questions are easier to answer before adoption than after.
- *Document AI use in conservation records.* When AI tools contribute to condition assessments, treatment planning, or other conservation decisions, that contribution should be recorded with the same specificity as any other methodological input. Transparency in documentation is a professional standard that applies regardless of the tool being used.

Before implementing its AI feature, Bev/Art released an AI policy and later, a full AI responsibility hub that outlines our commitment to applying AI with integrity, accountability and care.

[Click here to read more](#)



Looking Ahead

The trajectory of AI in the cultural heritage sector is toward broader integration and more sophisticated application. Shared standards, developed across institutions and in dialogue with affected communities, will be necessary to ensure that integration serves collections and the publics they belong to.

Conservators are well-positioned to lead in that conversation. The professional values that define conservation practice — material honesty, methodological rigor, long-term thinking, reversibility — are precisely the values that responsible AI adoption in cultural heritage requires. The field does not need to import an ethics framework from outside. It needs to apply the one it already has.

AI is a powerful interpreter of data but it is not a substitute for conservator judgment. Institutions that hold that distinction clearly, and build their adoption strategies around it, will be better equipped to use these tools well — and to recognize when they are being used poorly.