



COMMON EUROPEAN DATA SPACE FOR CULTURAL HERITAGE

Alignment paper on AI and the data space for cultural heritage:

*Boundary-setters and opportunity-seekers
working towards responsible AI futures*

Acknowledgements

We thank the Europeana community for their active engagement in the participatory process and their willingness to openly share their ideas, feelings and concerns about AI in the Alignment Assembly. We would also like to thank representatives from the Data space consortium, experts from the Europeana Advisory and Steering Boards, members of the Europeana Network Association and the Europeana Aggregators' Forum who attended the Alignment Assembly launch event and helped us create an initial set of statements that shaped our participatory process. We would like to express our gratitude to Alek Tarkowski, Director of Strategy at Open Future Foundation who provided guidance in setting up the Alignment Assembly.

This work has been carried out and funded under the contract with the European Commission (CNECT/LUX/2021/OP/0070: Deployment of a common European data space for cultural heritage). The [common European data space for cultural heritage](#) is a European Union flagship to accelerate the digital transformation of the cultural heritage sector. The data space increases participation in Europe's cultural heritage and promotes digital culture as a public good. It offers access to high-quality data, with an emphasis on 3D, and encourages people to reuse it and find new value in it.

Credits

This document was developed by the authors below, and is available under a Creative Commons CC BY-SA 4.0 licence. Cover image: [Hanna Barakat & Archival Images of AI + AIxDESIGN](#), [CC BY 4.0](#).

Emma O'Hare, Rasa Bočytė, Johan Oomen, Netherlands Institute for Sound & Vision, The Netherlands
Lorena Aldana, Antoine Isaac, Harry Verwayen, Valentine Charles, Europeana Foundation, The Netherlands

October 2025



Abstract

If the digital cultural heritage community is to play a meaningful role in the development of AI, it must first develop a shared vision. Yet the vague use of terms such as ethical, trustworthy and participatory AI, combined with the lack of concrete practices to make these concepts tangible, creates complexity and uncertainty. As the Europeana Initiative and its partners scale up the common European data space for cultural heritage, the question becomes urgent: **what role should the data space take in responsible AI development and adoption?**

This paper is the first outcome of the “[Alignment Assembly on Culture for AI](#)”, a collective intelligence exercise initiated within the data space, which engaged circa 400 professionals. It maps the results of the community consultation, identifies areas of consensus, friction and uncertainty and pinpoints topics that require further exploration. The paper identifies two contrasting opinion groups: boundary-setters (32% of respondents), critical of AI use, and opportunity-seekers (68%), advocating for its adoption. It explores key themes and insights that emerged from this exercise - ranging from defining and defending the sector’s position and responsibilities, to positioning heritage data as responsible AI training material, using AI in a way that is climate responsible or preventing the perpetuation of biases when AI is used with heritage data, among others.

Building on these insights, the paper outlines directions for further community work and proposes a set of priorities, an action plan, and immediate next steps. The Alignment Assembly thus serves as the starting point for a deeper, ongoing dialogue within the digital cultural heritage community about its engagement with and stance on AI technologies.

Table of contents

| | |
|--|----|
| Introduction | 5 |
| Our method: collective intelligence | 11 |
| Alignment Assembly results | 16 |
| Directions for further community work | 35 |
| From insights to action: priorities and next steps | 41 |

Introduction

Artificial Intelligence (AI) is certainly not a new topic in the cultural heritage community; institutions and professionals started developing and applying sector-specific machine learning techniques two decades ago. What has changed recently, and continues to change rapidly, is the scale, scope, pervasiveness and impact of AI techniques, as well as its societal impact. Today, AI is embedded in industrial processes, business models, as well as policy agenda and public discourse.

In parallel to these shifts, the Europeana community has been shaping the digital transformation of the cultural heritage sector, developing frameworks, standards and infrastructures that have guided the sector through key dilemmas, challenges and developments. **The digital cultural heritage community has never been a bystander in digital transition; we have actively shaped and developed digital infrastructures and tools that adhere to public values and inspire other sectors.** This ranges from developing technical standards such as the International Image Interoperability Framework (IIIF), to championing Creative Commons licences and data governance frameworks. While as a cultural heritage community we have made significant progress, much work remains to be done to fully realise our vision to democratise access to cultural heritage data, foster a culture of openness and boost Europe's creative capital. At the same time, we must continue to adapt to a rapidly evolving landscape and expand our relevance by sharing our expertise across other domains and driving greater social impact.

Existing use cases demonstrate the potential of AI tools to support the automation of labor-intensive tasks, improve the quality of heritage data, enhance storytelling around collections and open new ways of engaging with cultural heritage. However, the rapid rise of AI tools also brings considerable uncertainty. Professionals across the sector face important dilemmas around its ethical use, its impact on the trustworthiness of heritage data, and the future of human expertise. In some cases, AI systems champion the principles of openness and transparency, sharing and

fair reuse upon which the digital heritage community has been built, in many others, undermine or exploit them.

The common European data space for cultural heritage - a European Union flagship initiative to accelerate the digital transformation of the cultural heritage sector, with the Europeana Initiative at its heart - offers a unique opportunity to address these complex challenges and dilemmas.

If the digital cultural heritage community is to claim a meaningful role in the development of AI, we need a shared vision. However, the vague and ambiguous use of terms like ethical, trustworthy and participatory AI, combined with the lack of concrete good practices to make these concepts tangible, creates complexity and uncertainty

As the Europeana Initiative and its partners scale up the common European data space for cultural heritage - expanding its infrastructure and nurturing its community, the question becomes more pressing than ever: what role should the data space play in responsible AI development and adoption?

If we are to stay relevant in today's fast-changing world, the digital heritage community cannot stand by and simply observe from afar how AI is developed and applied across various societal domains. As mission-driven, public organisations, and as individuals with extensive expertise in digitisation and culture making, the community is well-positioned to advocate for more responsible and critical AI practices.

For this reason, the Europeana Foundation and the Netherlands Institute for Sound & Vision, on behalf of the data space consortium, launched in spring of 2025 an [Alignment Assembly on Culture for AI](#) - **a participatory collective intelligence initiative and consultation process designed to help our community clarify its position on the development, implementation and use of AI.**

This paper is the first result of the Alignment Assembly, which brought together circa 400 professionals. It maps the results of the community consultation, identifies areas of consensus, friction and uncertainty and pinpoints topics that need further exploration.

How to read this paper

This paper is written for and with the input of the digital cultural heritage community in Europe. Professionals, experts and volunteers working in, with and around cultural heritage, as well as data space stakeholders were consulted during this process. Additionally, the results of this work can provide valuable insights for policymakers and offer inspiration to other sectors that are also grappling with their stance on AI.

This paper starts by providing context on the common European data space for cultural heritage and its engagement with AI technologies so far, and its connection to the growing focus on AI at the EU policy level. Additionally, it provides clarification on what we mean—and do not mean—when we use the term ‘Artificial intelligence (AI).’ It then details the applied community consultation methodology and deliberative process, followed by an analysis of key themes and voting results, covering areas of consensus, friction and uncertainty. It concludes with some key take-aways, recommendations for further work and next steps.

The Alignment Assembly is the starting point for a deeper, ongoing dialogue within the digital cultural heritage community about our engagement with and stance on AI technologies. Given the rapid pace of AI advancements, this dialogue must be continuous and inclusive, involving as many stakeholders of the common European data space for cultural heritage as possible. It calls for active commitment to experimentation, exploration and critical analysis of AI use in the data space, while fostering community discussion and sharing of desirable and undesirable practices.

AI within the European data space for cultural heritage

AI is a broad term, and used in often unclear ways. In simple terms, AI refers to the use of computational methods to process information, identify regularities, and automate reasoning or decision-making tasks that would normally require human intelligence. Today, however, when people speak of AI, they often refer to a specific type: generative AI, or Large Language Models (LLM) - a type of machine learning model designed to generate content such as text images, speech or videos, based on training data. Yet machine learning is a far broader field of Artificial Intelligence focused on systems that can learn from data and thereby the systems can improve their performance over time without being explicitly programmed for specific tasks. Moreover, machine learning is also only a subset of AI. **Thus, when the term AI is used in this paper, we do not only refer to LLMs or even only machine learning, but the broad spectrum of AI technologies.** However, the focus will dominantly be on tools that are most relevant - and already used - to the cultural heritage sector, as exemplified in the next section.

Over the past decades, the Europeana Initiative has driven the digital transformation of Europe's cultural heritage sector. In recent years, as the Initiative has taken on stewardship of the common European data space for cultural heritage, AI has moved to the centre of its work. The data space comprises cutting-edge infrastructure, a vibrant community and a suite of products, frameworks and tools which facilitate the open and trustworthy sharing of heritage data across Europe.

The [European Commission Recommendation](#) on a common European data space for cultural heritage (2021) encourages the use of AI to improve digitisation and preservation processes, raise content quality and foster upskilling in the sector, positioning the data space as an active driver of AI adoption. Moreover, the core mission of the data space is to democratise access to cultural heritage data; making it findable, openly available, trusted and easy for anyone to use. AI can play a key role in accelerating this mission. Various AI experiments are already underway across the data space: from data enrichment to super-resolution (SR) or detecting and mitigating biases in heritage collections. These use cases are briefly described below.

Some examples of AI applications in the data space

Data enrichment. AI is currently being used for metadata enrichment on thousands of records in Europeana.eu that only have a title but lack a description. Data-to-text (D2T) methods convert structured information into natural language, creating missing metadata (such as dc:description or dc:title) in a clear, human-readable form. An AI-powered image enhancement algorithm is also being applied on low resolution thumbnails, automatically doubling resolution. Super-resolution techniques (SR) can unblur areas of interest and remove compression artifacts. This solution has already been deployed on more than one million images, giving users clearer previews and making cultural heritage more accessible and usable.

Detecting biases through AI-powered tools. Another impactful use of AI in the data space is for the detection of biases in heritage collections. Cultural heritage institutions have been describing their collections for decades, yet, this metadata is rarely updated to reflect changes in language and society. As a result, many descriptions that once aligned with dominant social narratives now convey outdated views, ignoring or alienating diverse communities and, in some cases, using language that is inappropriate, or even harmful. [DE-BIAS](#), a EU-funded data-space supporting project has analysed over 4.5 million records on Europeana.eu, automatically detecting problematic terms in heritage metadata in collaboration with affected communities. The DE-BIAS Vocabulary identifies almost 700 contentious terms used in cultural heritage metadata in five languages.

The policy landscape

AI is gaining attention at the EU policy level - the very environment in which the Europeana Initiative and the data space operate. Regulatory frameworks and funding mechanisms on AI are rapidly evolving and gaining momentum. With the adoption of its first legal framework on AI, the [AI Act](#), the EU is pioneering regulation to mitigate the threats posed by AI and algorithmic systems. This binding legislation establishes a uniform legal framework that classifies AI applications based on their risks to health, security, fundamental rights, the environment, democracy and the rule of law. On this basis, it prohibits applications deemed to pose unacceptable risks and imposes obligations on high-risk applications. Additionally, it mandates general transparency obligations for generative AI and the identification of 'deep fakes' as such.

The European Commission also aims to champion the innovative potential of AI, with the ambition to become 'the AI continent'. Access to data is at the heart of the AI Continent action plan to leverage the potential of AI and address its risks, with the aim of enhancing Europe's competitiveness and innovation, safeguarding its democratic values and protecting its cultural and linguistic diversity. **In this context, the data space for cultural heritage has huge potential to enable access to high-quality cultural heritage data for AI researchers, innovators and developers, promote the use of AI by cultural heritage institutions, and strengthen AI skills in the sector.** At the same time, it can become a place where the promises about societal, economic and democratic potentials of AI are put to a test and validated or refuted.

Against this backdrop, the Commission is currently developing sector-specific strategies to accelerate AI adoption in public services, including culture — through initiatives such as the Apply AI strategy and the forthcoming Cultural Compass. To ensure these AI frameworks truly serve the public interest, it is vital to embed cultural heritage perspectives — especially their commitment to openness, transparency and the long-term stewardship of authenticity in digital ecosystems. **This brings us to a central question: How can our sector leverage the data space to actively shape evolving EU policies?**

Our method: collective intelligence

The debates around AI have a tendency to be either highly dystopian or utopian. **In the Alignment Assembly, we adopt the position that we should treat AI as [‘normal’ technology](#)**, recognising its potential but refraining from thinking of it as a highly autonomous potentially superintelligent entity in the near future.

To explore the key questions and dilemmas, and to guide the consultation process, we examined AI in cultural heritage through the lens of the thinking model below, which outlines where we think our efforts can be most consequential in relation to the impact of AI on our work.

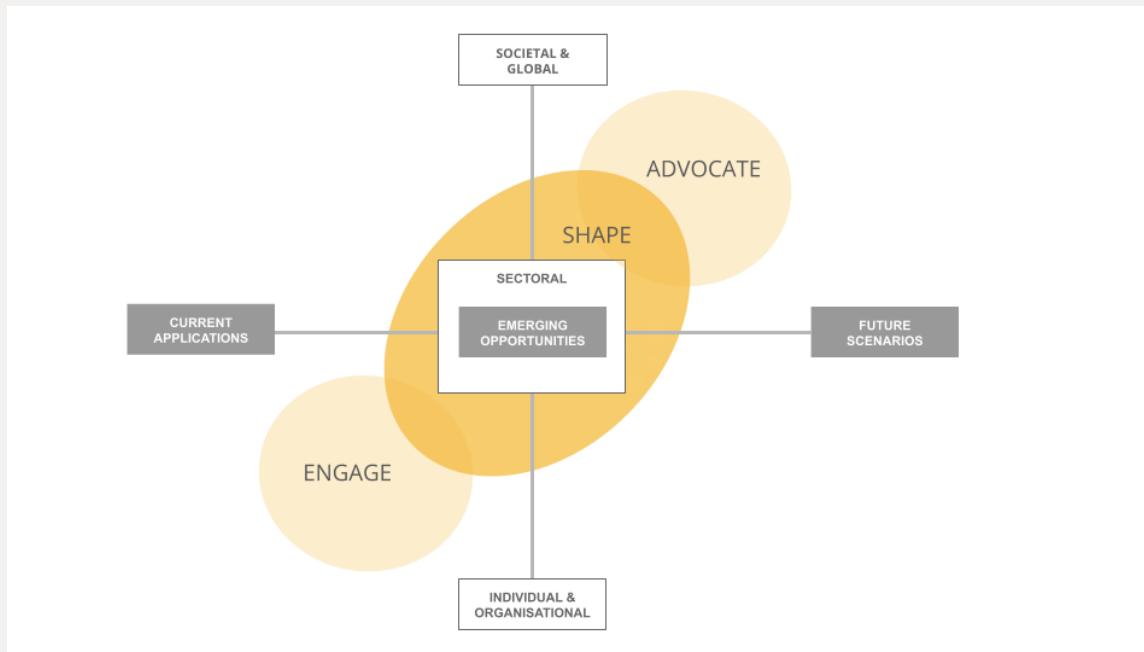


Figure 1: Engage/shape/advocate model

This model illustrates how we see AI development in relation to the impact it has on the data space, and consequently where we believe we need to focus our attention. The X axis represents the temporal dimension of AI deployment — from current applications, to emerging opportunities, to future scenarios. The Y axis captures the different groups involved and impacted: individuals and organisations, the heritage sector as a whole, and broader societal and global stakeholders.

The focus of our attention is primarily on the dilemmas that occur where AI can have sectoral impact through the widespread adoption of still emerging technology.

This is where we are in a position to **shape** what is being developed and under what conditions we can foresee widespread adoption. While doing so, we uncover a pattern of values that we can also apply to the tools and services we already **engage** with today but are not entirely sure about and it puts us in a position to **advocate** for the types of services that not only positively impact our sector but society as a whole.

The Alignment Assembly: a participatory exercise


In order to map and examine existing perspectives and practices in the digital heritage community, we chose to base the process on the Alignment Assembly model developed by the [Collective Intelligence Project](#) and inspired by the Alignment Assembly on “[AI and the Commons](#)” conducted by Open Future, Creative Commons and Fundación Karisma in February-March 2024. This approach combines large-scale surveys with structured conversations to inform policy debates and guide technology development in line with shared values.

[The Alignment Assembly model aims to “bring technology into alignment with collective values”](#). The survey took place on pol.is, an online open source survey platform described as: “a real-time system for gathering, analyzing and understanding what large groups of people think in their own words, enabled by advanced statistics and machine

learning.” The platform helps to map out opinion groups and show the areas of consensus, friction and uncertainty. Moreover, it also encourages participants to submit their own statements.

The starting point for this participatory, collective intelligence exercise was its online launch event in May 2025. Data space consortium partners, experts from the Europeana Advisory Board and Steering Board were invited, as well as members of the Europeana Network Association, Europeana Aggregators forum and experts from Netherlands Institute for Sound & Vision. During this first participatory session, an initial set of statements was developed, refined and reviewed. A first selection of 20 statements were published on the pol.is platform and the survey was publicly launched and widely distributed afterwards, including through Europeana website, relevant newsletters and

Welcome to a new kind of conversation — vote on other people's statements — the more the better.

 Anonymous wrote: 3 remaining

Cultural heritage institutions shouldn't use AI tools for metadata until they meet quality standards

☒ Agree ☐ Disagree ☐ Pass / Unsure

Figure 2: An example of voting process on the pol.is platform

Participants could anonymously cast their votes on published statements, and submit their own additional statements for the community to vote on until July 4th. The survey page on the pol.is platform was accompanied by clear guidelines to participate:

- State a clear position that people would feel compelled to agree or disagree with;
- Do not shy away from being intentionally provocative;
- Feel free to phrase statements that you might disagree with;
- Make statements that we do not know the answers to and want to explore with the community;
- Do not use jargon or terms that only AI specialists would be familiar with;
- Keep in mind that the scope of this mapping is AI in the cultural heritage field. Try to phrase statements that reflect concerns of this field.

Are your perspectives or experiences missing from the conversation? If so, add them in the box below — one at a time.

What makes for a good statement?

- *A stand-alone idea*
- *A new perspective, experience, or issue*
- *Clear & concise wording (limited to 140 characters)*

*Statements are displayed randomly and you are not replying directly to other people's statements: **you are adding a stand-alone statement.***



Share your perspective (you are not replying — submit a stand-alone statement)

Submit

Figure 3: Adding a new statement on the pol.is platform

Colleagues from the Europeana Foundation and The Netherlands Institute for Sound & Vision overseeing the Assembly carefully reviewed these submitted statements, according to the following criteria:

- They should fall within the scope of the topic of the Alignment Assembly as outlined in aforementioned engage/shape/advocate model;
- They should articulate a clear standpoint that would provoke voters to agree or disagree;
- They should be concrete and avoid vague terms that do not refer to concrete actions;
- They should avoid jargon to ensure that people with different levels of familiarity with AI could understand them;
- They should be distinctive enough compared to the already published statements.

Besides these criteria, we also wanted to avoid overwhelming participants with too many statements during the short voting period, as this might have led to fewer responses and potentially diluted the results.

The next chapter will dive into the voting results and prevalent themes that were distilled from the submitted statements. As stated earlier on, these findings are the starting point for further dialogue and exploration of dilemmas. This paper includes recommendations based on the results and some follow-up events and collaborations are already lined up (see Next Steps at the end of this paper).

Alignment Assembly results

The Alignment Assembly was open for votes and statements on the pol.is platform between 26 May 2025 and 4 July 2025. In total, 372 participants cast their votes and a total of 10.242 votes were cast over 54 statements.¹ The survey was initially launched with a set of 20 statements derived from the online Alignment Assembly launch event on 15 May 2025.

Following the launch, an additional set of 103 statements was submitted by participants², out of which 34 were selected for public voting. Some of the statements submitted by the community lacked clarity or repeated already covered points. Some statements submitted and published towards the end of the voting period did not get as much engagement, and for this reason may not provide convincing results. However, we did analyse all statements that were submitted in order to determine key themes and topics around AI and cultural heritage.

The voting outcomes can be grouped into three categories, each described in the following section.

Three categories: consensus, friction and uncertainty

The analysis of the Alignment Assembly results shows that community perspectives on AI in cultural heritage cluster into three main categories: areas of strong consensus, points of friction, and zones of uncertainty. These categories provide a useful framework for understanding where the community already shares clear positions, where significant divisions remain, and where further exploration and discussion are needed.

¹ The dataset with all the statements is openly available on Europeana's Zenodo account: <https://doi.org/10.5281/zenodo.17252598>

² Excluding duplicate statements.

Strong consensus. Statements where the community aligns to a large extent (with a 60% or higher percentage of agreement¹); 23 of the 53 published statements show a high level of consensus among the voters. Consensus can be found across several themes as we shall lay out in the next section.

Clear divide. Alongside areas of consensus, the results also reveal tensions around statements that have no overwhelming agreement or disagreement overall. Two opposing opinion groups have materialised from the pol.is results:

- Group A, approximately 32% of the voters, or 93 participants, are clearly very critical of AI use in the cultural heritage sector. From this point onwards, we will call them the **boundary-setters**.
- Group B, approximately 68% or 195 participants, advocates for the importance of embracing AI tools in the cultural heritage sector: the **opportunity-seekers**.

A total of 15 statements with lower general levels of (dis)agreement demonstrate this divide by showing greater levels of (dis)agreement between these two groups.

Areas of uncertainty. Besides consensus and friction, the reception of some statements shows uncertainty, as a large percentage of voters neither agreed nor disagreed with a statement, but chose to click on pass/unsure instead. They remain open questions.

While some participants might have chosen to skip a statement because they lacked sufficient knowledge to express an opinion or the language was not clear, in other cases it might be that voters found it hard to imagine how the presented statement would play out in real life or they might have mixed feelings about it.

¹ Some of the statements that were included later in the process, received fewer votes. We realise these results are less reliable. Therefore we focus mainly on statements with a significant number of votes.

For instance, in the case of using AI summary features in editorial processes (statements #90 in the Figure 4 below), the lack of existing case studies that dive into pros and cons might be a factor behind the uncertainty. Similarly, Figure 4 statement #55 presents a scenario of using principles similar to the fairtrade mark being applied to AI tools; thinking exercises that would develop this idea further (e.g. how would the ‘fairness’ of tools be determined and who would do it? How would the costs differ from the current offer? How would that impact access across the sector?) would help to bring more clarity about the community’s position on them. In this sense, areas of uncertainty provide a promising direction for ideation activities with “what if” scenarios.

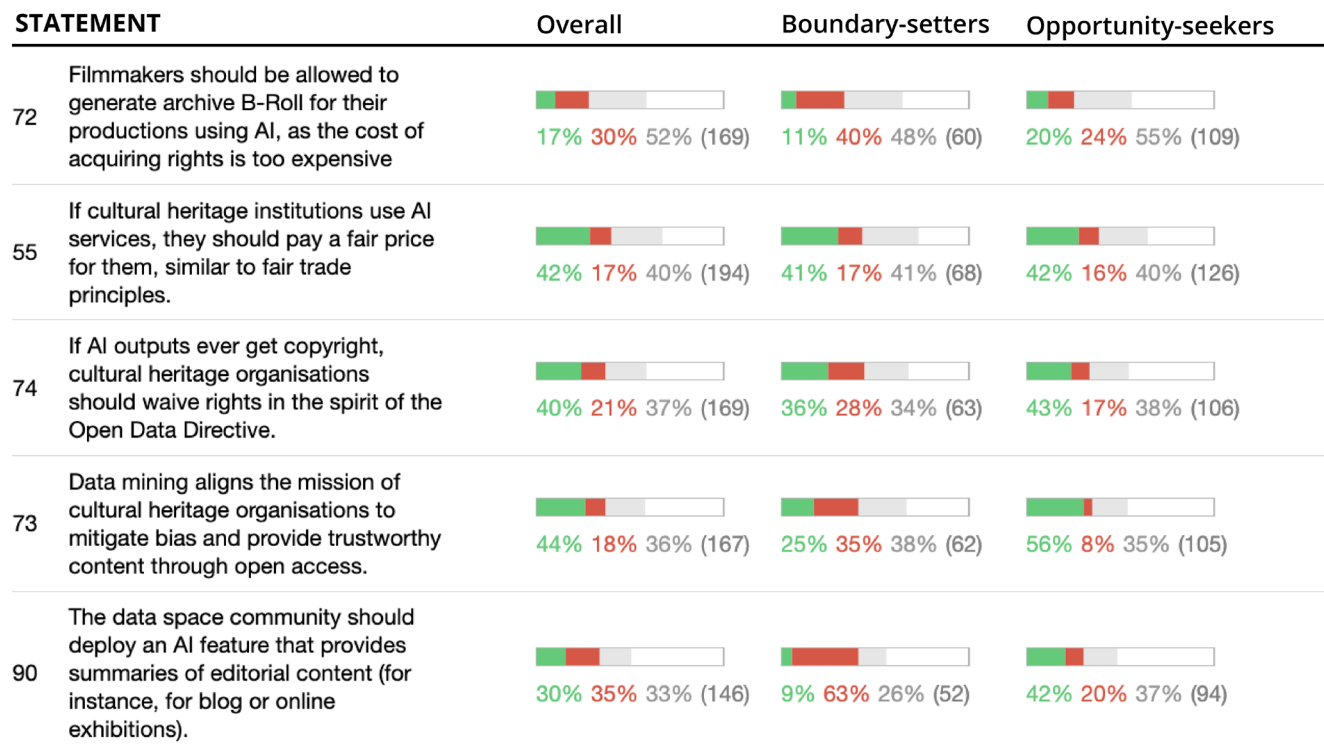


Figure 4: A snapshot of statement

Key themes and insights

This paper analyses all 123 submitted statements, including the ones that were not published for voting. With the thinking model of Figure 1 as a framework in mind, **seven themes** that the sector wants and needs to position itself on emerge, raising questions that demand further bottom-up and top-down discussions and action. In this chapter we will unpack the themes and connect them to the voting results. As we develop the data space, it will be important to address them:

- 1. Defining our position and responsibilities as the foundation for shaping policy and advocacy:** Where should the data space for cultural heritage position itself in relation to concerns and responsibilities boundary-setting and opportunity-seeking signals from the heritage community? What policy-influencing, advocacy and narrative shaping positions should the data space community take on?
- 2. Positioning heritage data as responsible AI training material.** Should we build infrastructures that actively encourage, deter, or explicitly prohibit the creation of training datasets for AI applications? Who gets to use our data, and how does that challenge our stance on open access? How can data authenticity and integrity be ensured? How do we reconcile the principle of open access with the need to control who uses our data and for what purpose?
- 3. Taking climate action seriously.** Can we keep ourselves and others accountable for the environmental impact of AI deployment in the data space? How do we deal with the lack of transparency around the topic?
- 4. Preventing the perpetuation of biases in AI development.** Will the use of our data in AI applications reinforce systemic injustices? How can we build more inclusive and fair value chains around heritage data across diverse languages and cultures?

- 5. Building capacity and developing AI literacy in the sector.** What workforce transitions should we prepare the community for? And what skills are essential for heritage professionals to confidently understand and assess, and use or reject AI tools?
- 6. Forging careful alliances beyond the heritage sector.** Can we embrace collaborations on AI topics with other sectors while safeguarding our values? What rules of engagement need to be put in place to ensure trust, fairness and shared benefit?
- 7. Clarifying what makes an AI use case desirable or undesirable.** Should the data space take a stance and encourage certain AI deployment use cases while discouraging others? What principles can help us identify responsible applications to promote while remaining flexible and open to embrace future scenarios?

It is worth noting that given the grassroots nature of the Assembly, not all topics that are being discussed in the wider sector feature prominently in the consultation (take, for instance, questions of authenticity, community authority, or the long-term sustainability of infrastructures supporting AI use). That is not to say that these other topics are not important; the results should be interpreted as a snapshot of ideas that were dominant during the weeks the Assembly took place.

Theme 1: Defining our position and responsibilities as the foundation for shaping policy and advocacy

An overarching and very prevalent theme that can be distilled from the submitted statements is in regard to the responsibilities the voters attribute to the cultural heritage sector. While many of the other themes can be understood in relation to this overarching perspective, we consider it important to address this topic separately as well. It signals a strong desire and need for a sector-specific approach, as well as for the sector to actively claim agency regarding AI adoption based on the community's values. The statement with the highest level of consensus, with a 91% agreement rate, exemplifies this:

“The cultural heritage sector should develop its own narrative about AI, ones that are grounded in our values rather than focused on efficiency and commercial gains.”

(91% agree, 3% disagree, 5% pass)



A large majority of voters agree that AI technologies must align with the community values. Another statement that shows clear consensus also concerns the responsibilities of the heritage institutions around setting boundaries and protecting the sector:

“In the name of openness, cultural heritage data is being silently harvested to train commercial AI models —often without consent, credit, or compensation. Clearer boundaries and shared licensing models are needed to protect public assets from exploitation.”

(83% agree, 10% disagree, 6% pass)



A similar message is echoed in other consensus statements. In other words: the community should not just simply stand by but wants to be actively involved in setting the boundaries, conditions and narratives around the use of AI tools. Advocating for regulations that are tailor-made for cultural heritage is also part of this stance:

“The data space for cultural heritage should advocate for sector-specific AI regulations at EU level.”

(70% agree, 10% disagree, 19% pass)



Areas of tension between the boundary setters and opportunity seekers also surface in the voting results regarding what position and responsibilities the sector should and should not embrace, as shown in the Figure 5 below.

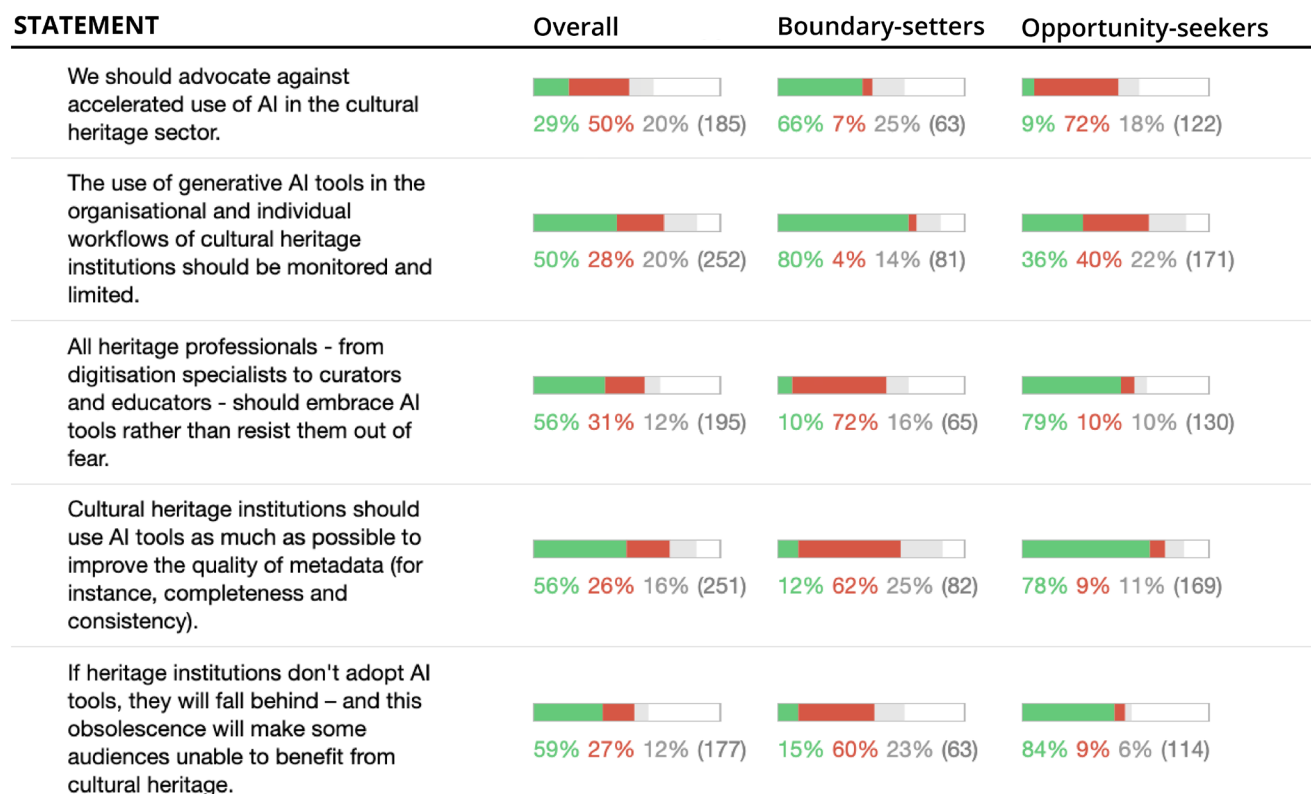



Figure 5: Non-consensual statements that hint at the existence of two opposite groups


These dividing statements show quite a fundamental difference of opinion regarding the positioning of the sector. The boundary-setters are distinctly resistant towards accelerating the focus on AI tools in the sector and argue that their use should be monitored and limited. On the other side, the opportunity-seekers consider it a necessity to embrace AI tools as much as possible and fear that the heritage institutions will fall behind if they do not. There is friction and uncertainty on what should be the core values of heritage institutions:

These dividing statements show quite a fundamental difference of opinion regarding the positioning of the sector. The boundary-setters are distinctly resistant towards accelerating the focus on AI tools in the sector and argue that their use should be monitored and limited. On the other side, the opportunity-seekers consider it a necessity to embrace AI tools as much as possible and fear that the heritage institutions will fall behind if they do not. There is friction and uncertainty on what should be the core values of heritage institutions:

“Being AI-free and human-centered should be a fundamental, core value of heritage institutions.”



“The common European data space for cultural heritage should abandon AI efforts instead focus on community outreach, editorial work and human enrichment of data.”




Although a majority of the voters disagree with the statements above, it does reveal tensions between the boundary setters and the opportunity seekers. These areas of friction require further dialogue, as they are rather elemental to determining the sector’s position on AI.

Theme 2: Positioning heritage data as responsible AI training material


Naturally, the topic of data access is an important theme. Heritage institutions manage a lot of data, AI companies need data to train and improve their models. The opinions captured in the Assembly offer a mixed picture as out-looks differ, and there is considerable uncertainty around this topic.

The core challenge is that there is broad support in the heritage sector that publicly funded heritage data should be made as openly available as possible to support reuse, research and innovation. On the other hand, privacy concerns, rights management, and the potential use of opt-outs by institutions or rights holders risk limiting access, sometimes even to public-domain materials. In the Assembly, this issue was addressed in the following dilemmas, where the level of agreement is quite high but the level of disagreement is significant too:

“Cultural heritage institutions should keep their data as open as possible, even if AI models trained on it generate commercial profit.”



“The best way for cultural heritage institutions to ensure cultural heritage is well represented in AI is by making cultural heritage data available.”



The boundary-setters are divided on this issue, whilst the other group, by majority, advocates for openness, even if AI companies generate profit from it.

When it comes to the opportunity-seekers (68% of participants), this group prefers keeping data open, whilst the boundary-setters advise cautiousness. Other statement voting results do tend to paint a more nuanced picture. Over 50% of both the boundary-setters as well as the opportunity-seekers believe that:

“The common European data space for cultural heritage should curate datasets for the training of AI systems only if these are open-source solutions, developed within the EU.”

(54% agree, 17% disagree, 27% pass)



However, 27% passed on this statement. It points to some uncertainty around the topic. Further discussion might be needed. Does this point to some doubt whether datasets should be shared for AI-training purposes at all, or does it point to the thought that the AI training systems do not have to necessarily be exclusively open source? A larger majority agrees on where the focus of the institutions should be regarding their data:

“Rather than limiting the use of heritage data for training AI models, the focus of cultural heritage institutions should be on ensuring proper provenance and attribution of heritage data.”

(72% agree, 11% disagree, 16% pass)



Even 53% of the boundary-setters support this statement. Also, most voters from both groups agree that investing in open-source alternatives is essential:

“Investing in open-source and interoperable AI alternatives is essential to uphold the public values of mission-driven institutions.”

(76% agree, 8% disagree, 14% pass)



Common ground can definitely be found regarding data access, but a more thorough understanding of the different viewpoints is needed in order to determine the community's position.

Theme 3: Taking climate action seriously

Setting standards and boundaries does not only relate to dealing with data. Retaining agency is also considered crucial when considering climate impact. There is a large consensus among the voters regarding this theme:

“Environmental responsibility standards should be developed by the cultural heritage sector as a requirement for the use of AI tools in the data space.”

(75% agree, 12% disagree, 12% pass)



“There must be serious attention given to the environmental impact of digital culture and AI before more infrastructure is built. Slow down!”

(71% agree, 18% disagree, 10% pass)



“Heritage organisations should only use generative AI tools from companies that have concrete ecological strategies (for instance, planting trees to neutralise carbon footprint).”

(53% agree, 22% disagree, 24% pass)



The above statements clearly summarise why this topic should play a part in the sector’s stance on AI. Interestingly, none of the boundary-setters disagreed with the first statement, whilst still 28% of the opportunity-seekers voters disagreed. There is some imbalance between the two groups, but overall, a majority agrees that climate impact is a factor to be taken seriously; the second statement reaffirms this.

Theme 4: Preventing the perpetuation of biases in AI development

The community takes a convincing stance on bias as well. Voters agree that data contains bias, and we should therefore be cautious and alert about this:

"AI should be used alongside human expertise to prevent historical biases and misclassification."

(83% agree, 5% disagree, 11% pass)



"AI will have biases if trained on cultural data (on the data of non-Western cultural heritage held in European cultural institutions)."

(60% agree, 15% disagree, 24% pass)



"EU heritage institutions should collaborate with non-EU ones to ensure diversity, no bias, and mutually beneficial exchanges."

(85% agree, 6% disagree, 8% pass)



Quite a few statements that were submitted concerned the topic of (cultural) bias and diversity. Many submissions were about bias or related themes such as culture, diversity, equity and fairness. It proves to be a strong value in the heritage community and should be taken seriously moving forward.

Theme 5: Building capacity and developing AI literacy in the sector

A majority of the voting community agrees that increased focus on AI literacy and education and training of the heritage community is needed. Here, AI literacy is not limited to technical know-how but includes understanding the impact on jobs in the sector, as well as the cultural and societal shifts happening due to the mainstreaming of AI. Several statements confirm this outlook:

“The data space for cultural heritage should engage less in the technical development of AI tools and focus more on AI literacy and advocacy activities.”

(60% agree, 18% disagree, 20% pass)



“Next to investing in AI literacy of heritage professionals, public-facing heritage institutions also have a responsibility to educate the public about the importance of responsible use of AI.”

(82% agree, 9% disagree, 8% pass)



“The impact of the use of AI tools on skill development of different job roles in the sector should be studied.”

(87% agree, 2% disagree, 10% pass)



“Cultural heritage organisations should focus less on the technical capabilities of AI and instead address cultural and societal issues (like education, digital literacy) around it.”

(59% agree, 21% disagree, 19% pass)



There is some concurrence, although to a lesser degree, that heritage professionals currently lack knowledge or the capabilities to cast verdicts on AI use in the sector. This statement might, to some extent, explain why AI literacy is significant to a large number of voters:

"In general, cultural heritage professionals do not have enough understanding of AI to make decisions on its development."

(54% agree, 26% disagree, 18% pass)



Besides the voting results, other submitted statements that did not end up being published, also emphasise the importance of this subject matter:

"The data space for cultural heritage should offer AI training for heritage professionals."

"Development of higher skilled workers in cultural heritage will create new "employment" opportunities for AI"

"The heritage community should immediately begin developing AI training programs for heritage professionals."

Further community dialogue is needed to gather input on what kind of training is needed or wanted. There is for instance, disagreement about the following statement:

"Every cultural heritage professional should have basic coding skills and understand the logic behind AI systems."

(44% agree, 42% disagree, 13% pass)



Should the community be technically trained, or are there other areas that should be a priority? Further community dialogue is needed to assess which skills are needed in the heritage community.

Theme 6: Forging careful alliances beyond the heritage sector

Most voters remain cautious about forming partnerships with major tech corporations:

“Heritage organisations should reject partnerships with AI companies that fail to comply with ethical labour and data practices globally.”

(87% agree, 6% disagree, 6% pass)



“Heritage organisations should actively develop open-source AI models to provide alternatives to big tech’s control over cultural data.”

(65% agree, 17% disagree, 17% pass)



Once again expressing the view that the digital cultural heritage sector should act autonomously by setting its own (ethical) standards and even developing its own AI models:

“Cultural institutions should develop models that guarantee the veracity of the information offered by AI, and avoid falsehoods and lies.”

(78% agree, 4% disagree, 16% pass)



This does not mean that the community is against alliances or collaboration. However, partners should be aligned on their values. Several statement results underline this:

“The data space should actively work with LLM partners that share our values and accept that others use the data.”

(74% agree, 9% disagree, 16% pass)



“The data space for cultural heritage should collaborate with platforms like Hugging Face to develop standards for creating AI training datasets.”

(67% agree, 6% disagree, 25% pass)



“EU heritage institutions should collaborate with non-EU ones to ensure diversity, no bias, and mutually beneficial exchanges.”

(85% agree, 6% disagree, 8% pass)



“Cultural heritage organizations should stand up for the creation of European language models and cooperate with European universities.”

(84% agree, 4% disagree, 12% pass)



Clearly, voters in a large majority recognise the need for collaboration with European organisations as well as with organisations worldwide that prioritise diversity and the detection of biases. Partnering with largely open source initiatives, such as Hugging Face, points to valuing democratisation of machine learning.

Voters agree that, as long as the values align, alliances are necessary for researching, developing and and/or implementing AI tools. Not only do the published statements’ results point to this. There are also several unpublished submissions that concern partnerships and the necessary preconditions for collaboration:


“Trustworthy AI-supported digital transformation can only be achieved with cross-sectoral international collaboration.”

“Research on AI e.g. European LLM’s, should also happen in cooperation with applied science universities, not only universities.”

Theme 7: Clarifying what makes an AI use case desirable or undesirable


AI is an umbrella term for many different types of tools and technologies. Often when speaking of AI, people mean generative AI and refer to Large Language Models (LLM’s). But the scale and scope of AI is much larger and diverse and there are many possible use cases for the sector. Quite a few of the submitted statements concern thoughts on which tasks AI should perform and prioritise to improve the sector’s work. Of the statements that were published, disagreement on AI deployment becomes visible.

“Cultural heritage institutions should use AI tools as much as possible to improve the quality of metadata (for instance, completeness and consistency).”




Unsurprisingly, the boundary-setters (group A) do not support this view. The same goes for editorial content:

“AI techniques should not be used in content curation for editorial and education materials based on heritage data.”



“The data space community should deploy an AI feature that provides summaries of editorial content (for instance, for blog or online exhibitions).”




Not only contention, but also uncertainty regarding AI deployment becomes visible in Figure 12.

A couple of consensus statements point to some relativism about the deployment of AI in our sector. A majority agree that AI is not the future be-all and end-all. Neither running after AI alone nor supporting highly dystopian views is supported, which points to a healthy dose of realism about the impact of AI in the community. Two statements with a majority disagreement point to this:


“Cultural heritage institutions should drop their efforts to build visually attractive search interfaces and focus instead on implementing AI-driven services.”

(13% agree, 67% disagree, 19% pass)



“If AI tools become more popular, people will stop visiting libraries and museums, and only machines (like search engines and AI agents) will access collections.”

(13% agree, 74% disagree, 11% pass)



Most respondents do agree that AI can benefit organisations of different sizes, but this is to a lesser extent:

“Adoption of AI is so complex and costly that only large cultural heritage institutions will benefit from it.”

(21% agree, 63% disagree, 14% pass)



Since this is an area where the data space could have a direct impact, it will be worth diving deeper into the reasoning behind different responses to understand what is needed to provide adequate support for large, medium and small organisations.

Many other statements that were submitted, but unpublished, offered perspectives about what kind of AI deployment could benefit the sector. Stating, for instance, that focus should be on developing rule-based AI for automating the handling of digital material, not on generative AI tools. Another suggests AI should be used for easier search and research. While alternative submissions state that AI will never fully replace human restoration and preservation of digital cultural heritage. Yet another one states that the common European data space for cultural heritage should work on AI services, making use of collections from multiple institutions.

AI use cases demand further exploration. In which areas can AI tools benefit the sector, and under what conditions can we as a community get on board with them? How can the data space take a stance and encourage certain AI deployment use cases while discouraging others? Can we identify responsible applications to promote while remaining flexible and open to embrace future scenarios?

Directions for further community work

The high level of engagement on the [pol.is](#) platform shows the community is eager to collectively shape the role of AI in the data space for cultural heritage. With the focus brought by our thinking model, the Alignment Assembly was foreseen as a conversation starter that would lead to further community activities shaping the data space. As technologies develop at a rapid pace, it demands a flexible and adaptive approach.

Although the analysed data is insufficient to draw definitive conclusions from, it gives us a strong indication on where the community stands. It pinpoints dilemmas and issues that demand further exploration, but also consensus points where concrete actions can already be taken. It signals that the digital cultural heritage community must take action by crafting sector-specific narratives, setting boundaries, agreeing on shared values and investing in AI literacy. Some of those shared values have become substantiated: fighting bias, considering climate impact and partnering with external parties who share the same values. Other topics show more uncertainty or friction. While there are more consensual statements than divisive ones, our results indicate friction points regarding some crucial themes. Although the boundary-setters represent a minority in the voting results, it is not an insignificant group and therefore should not be ignored.

For each of the key themes and insights, we summarise consensus and friction points, and propose areas for further investigation below.

High consensus areas:

- Instead of echoing market-driven language about AI that focuses on efficiency and growth at all costs, the data space should develop its own narratives that represent the values of its community.
- The data space community can act as an advocate for sector-specific AI solutions, infrastructures and policies.
- AI should not be seen as an existential threat nor a utopian solution for digital heritage collections.

Friction points:

- Opinions differ on the extent to which the data space should promote and advocate the adoption of AI-supported workflows. Some see it as essential to accelerate efficiency and sectoral relevance, while others caution against pushing adoption without first addressing ethical, environmental, and social concerns.
- Calls for speeding up adoption as well as slowing down AI development coexist in the community, highlighting a clear dual sentiment amongst participants.

Potential avenues for further work:

- Develop a distinctive vocabulary and visuals that help to communicate the community's stance on AI. Use this in advocacy efforts within the sector as well as when positioning the heritage community within the broader AI landscape. The community stance should be rooted in ethical principles such as transparency, accountability and inclusivity, and align with Europeana's values: Usable, Mutual and Reliable.
 - Carry out further community consultation processes that provide deeper insights behind convictions and reservations about AI.
-

Positioning
heritage data as
responsible AI
training material

High consensus areas:

- The growing interest in heritage data as AI training datasets puts pressure to improve data quality and assure provenance.

Friction points:

- There are diverging opinions on whether the growing interest in using cultural data for training AI models asks us to revisit our commitment to openness.
- We note signals of disagreement on whether the data space should proactively contribute to the development of training datasets.

Potential avenues for further work:

- Monitor the use of the data space as a way to access heritage datasets for AI training.
- Investigate sustainable approaches to funding and maintaining open infrastructures as AI-driven data scraping increases their technical and operational demands.
- Develop community-led approaches to licensing, consent, and benefit-sharing that ensure transparent, fair, and accountable reuse of digital heritage in AI development.

Taking climate
action seriously

High consensus areas:

- The environmental impact should not be an afterthought in AI adoption but something meaningfully and practically considered from the very start.
- The scaling of AI implementation ambitions in the data space community should always factor in our environmental responsibilities.

Friction points:

- There were no disagreements on the importance of this topic, reflecting the community's clear commitment to environmental action. However, community statements remain vague on how to translate this into concrete measures.

Potential avenues for further work:

- Set up spaces for dialogue that bring together dispersed communities (climate activists, AI & heritage specialists, third-party technology developers, NGOs). whose views need to be represented in discussions about AI strategy.
 - Collaboratively develop and adopt frameworks to holistically measure the impact, including standardised metrics for energy, water, and carbon impacts.
 - Develop ways to communicate the environmental impact of AI technologies in easily understandable ways; identify the trade-offs in mitigating environmental impact.
 - Encourage debates about the pace and scale of AI adoption in the sector.
-

Preventing the perpetuation of biases in AI development

High consensus areas:

- While bias in data is inevitable, potential harms are amplified when heritage data is used as AI training datasets. This calls for careful processes that centre human expertise.

Friction points:

- There were no statements on this topic that indicated disagreement. However, while the potential problems are clear, community statements remain relatively vague on the practical measures for addressing them.

Potential avenues for further work:

- Promote initiatives that highlight harmful biases in cultural data and continue to spread awareness about this in the heritage community and beyond (including other data spaces).
 - Support activities that use cultural data to challenge current power imbalances and social injustices.
-

**Building capacity
and developing
AI literacy in the
sector**

High consensus areas:

- The data space should engage not only in technical development of AI tools but also to dedicate an equal amount of focus to workforce upskilling.
- The heritage sector has a societal responsibility to educate their audiences about AI and provide a nuanced perspective on the impact of these technologies on different areas of life.

Friction points:

- There is no clear consensus on what types of skills and knowledge are essential for cultural heritage practitioners to make decisions about AI. It is clear, however, that this is not limited only to technical skills.

Potential avenues for further work:

- Develop speculative scenarios imagining how jobs in the heritage sector might change with an introduction of new AI-driven processes.
- Investigate essential skills and knowledge that professionals across varying roles need in order to make decisions about AI (e.g. at policy and strategy level, in marketing and communication teams, etc.).

**Forging careful
alliances beyond
the heritage sector**

High consensus areas:

- Collaboration with cross-sectoral partners (from researchers and technology developers to standards setters) from diverse cultural contexts (within and outside of Europe) enriches the community with different perspectives and leads to mutually beneficial exchanges.
- We need to set rules for collaboration to ensure that our values are respected and cultural data (and the communities behind it) are not exploited.

Friction points:

- There were no strong disagreements on this topic and most statements encouraged the data space community to take a clear position on the types of collaborations it should encourage and reject.

Potential avenues for further work:

- Develop a list of principles to guide organisations when choosing who to collaborate with and under what conditions (including principles on what kind of partnerships are not desirable).
- Create case studies that showcase successful collaboration examples, as well as those that point to potential pitfalls that might result from clashing values.
- Identify priority areas where the data space community has unique expertise (e.g. setting standards for the descriptions of training datasets, copyright); proactively seek out cross-sectoral partnerships where we can share and build on this expertise.

Clarifying what makes an AI use case desirable or undesirable

High consensus areas:

- AI deployment in the community should not grow to such a scale where it overshadows other priority areas.
- AI deployment should not grow to such a scale where it becomes impossible to provide sufficient human oversight.

Friction points:

- There are differing opinions on the work processes (including data enrichment, curatorial and editorial work) which could be enhanced with AI and where such tools should be resisted.

Potential avenues for further work:

- Explore governance principles such as human-in-the-loop oversight, limits on automation of core archival work processes with an eye on safeguarding data authenticity.
 - Develop (speculative) case studies that unpack impacts of AI on different work processes, especially highlighting potential risks and impacts of prolonged use.
-

Priorities and next steps

Building on the key takeaways and directions for further work introduced in the previous sections, this chapter proposes a set of priorities and an action plan for how to build on those insights.

Developing ethical and strategic principles

Building on the Alignment Assembly insights, our priority is to shape a set of ethical and strategic principles that will guide the use of AI in the data space community. These principles will articulate our collective position and must remain true to Europeana's core values: Usable, Mutual, Reliable. We intend to make these principles emerge through a four-pronged collaborative process outlined below, which will seek to answer questions about the consensual priorities and open dilemmas.

1. Clarifying what makes an AI use case desirable or undesirable

How can the data space take a stance and encourage certain AI deployment use cases while discouraging others? Can we identify responsible applications to promote while remaining flexible and open to embrace future scenarios?

2. Unpacking critical dilemmas that emerged from the Alignment Assembly

Can we position heritage data as responsible AI training material? Should we design infrastructures that enable the use of heritage data for AI training, or should we deter or even prohibit it? How do we reconcile the principle of open access with the need to control who uses our data and for what purpose? Can we guarantee the authenticity and integrity of heritage data once it becomes part of AI training pipelines?

Can we use AI in a way that is climate responsible? How can we hold ourselves and our partners accountable for the environmental costs of AI in the data space? Do we risk complicity if we continue deploying AI without transparency around its energy and resource impact?

Can we prevent biases and be more inclusive when AI is used with heritage data? Will the use of our data in AI applications reinforce systemic biases and injustices, or can it help dismantle them? While heritage collections often reflect outdated or even derogatory perspectives of the world, can we ensure that heritage data contributes to inclusive, fair value chains that respect diverse languages and cultures, or at the very least are more transparent with respect to their biases? Can we promote best practices in this respect?

Can we forge ethical alliances beyond the heritage sector? Can we collaborate with other sectors on AI while still safeguarding our core values? What rules of engagement are necessary to build partnerships that ensure trust, fairness, and shared benefits?

3. Building capacity, developing AI literacy in the sector

What skills are essential for heritage professionals to confidently understand and assess, and use or reject AI tools? And what workforce transitions should we prepare the community for?

4. Advocating, defending our position, influencing policy

Where should the data space for cultural heritage position itself in relation to boundary-setting and opportunity-seeking signals from the heritage community? What policy-influencing, advocacy and narrative shaping positions should the data space community take on?

Immediate next steps

The initial analysis outlined in this Paper provides valuable insights into the community's position on several key AI-related topics, highlighting areas of consensus, divide, as well as questions that require further exploration. This is only the beginning. The work will continue through a series of planned actions within the common European data space for cultural heritage throughout the next year (2025-2026):

1) Joining forces with like-minded organisations and networks (ongoing): An ongoing collaboration with AI4LAM — a network committed to transparent, human-centric AI — is already under way to jointly accelerate and advocate for responsible AI in cultural heritage. The Europeana Foundation, as leader of the data space consortium, will formally join this network to ensure knowledge exchange, and to keep the data space at the forefront of global developments. The key insights of this paper will be presented at the Fantastic Futures conference by AI4LAM, taking place on 3–5 December 2025 at the British Library in London. Similarly, we are exploring collaboration with the Institutional Data Initiative at Harvard Law School Library and the Open Future Foundation.

2) Continuing community dialogue (ongoing): Building on the momentum created by the Alignment Assembly, and as an effort to deepen and continue this dialogue, the EuropeanaTech community of the Europeana Network Association and AI4LAM will organise a cross-community workshop and webinar on AI, taking place in Copenhagen on 27–28 October. The workshop will review emerging insights, incorporate perspectives from diverse domains of expertise, and contribute to the development of a sectoral white paper on responsible AI use in cultural heritage.

3) Advocating to policymakers with the insights from the Alignment Assembly (ongoing): The data space can play a meaningful role in shaping ongoing AI policy discussions at the EU level. Findings from the Alignment Assembly analysis will help define sectoral positions and inform policymakers.

