

THE RESPONSIBLE AI AFTERLIVES WORKBOOK

Exploring AI to 'revive' collections and engage visitors



Jenny Kidd, Eva Nieto McAvoy, Bethan Jones and Alison John



LEVERHULME
TRUST



Arts and
Humanities
Research Council



Contents

- 4-5 – Setting the scene: what do we mean by 'AI afterlives'?
- 6-7 – Introducing the Responsible AI Afterlives Workbook
- 8-10 – **Part 1 – Values first**
- 11-12 – **Part 2 – Why AI afterlives**
- 13-14 – **Part 3 – Why not AI afterlives**
- 15-18 – **Part 4 – Ethics & responsible AI**
- 19-23 – **Part 5 – What to consider when designing responsible AI afterlives**
 - Creativity & care
 - Community impact
 - Accessibility & inclusivity
 - Data integrity
 - Reflection & iteration
- 24-27 – **Part 6 – Your AI afterlives action plan**
- 28-29 – Glossary of key terms

This workbook has been produced with support from the Leverhulme Trust, the Arts and Humanities Research Council, Cardiff University and King's College London. It is based on our Synthetic Pasts project, which investigates how AI and automation are being used to "revive" archival materials.

Find out more or get in touch at www.syntheticpasts.com

Special thanks to the cultural professionals who participated in workshops and provided feedback throughout development.

Researchers: Bethan Jones, Jenny Kidd, Eva Nieto McAvoy

Project Partner: Alison John, yello brick <https://www.yellobrick.co.uk>

Research Assistance: Rio Creech-Nowagiel

If you use the guide in research or professional practice, please cite it as follows:

Kidd, J., Nieto McAvoy, E., Jones, B. and John, A. (2025). The Responsible AI Afterlives Workbook: exploring AI to 'revive' museum collections and engage users. Synthetic Pasts. <https://doi.org/10.5281/zenodo.17619346>

© 2025 Synthetic Pasts

Designed by yello brick

Licensed under CC BY 4.0 (creativecommons.org/licenses/by/4.0/).

This licence permits sharing, adaptation and reuse with appropriate attribution.

Published by Synthetic Pasts

DOI: 10.5281/zenodo.17619346


Released November 2025



www.syntheticpasts.com

Image credit for cover and logo: Google's

Duet AI, 30 Nov 2023, prompt by JK



Setting the scene: what do we mean by 'AI afterlives'

Artificial Intelligence (AI) is reshaping the way cultural heritage organisations work with collections and engage with visitors. In time, it might even impact the stories they tell. From automated cataloguing to immersive visitor experiences the possibilities are exciting, but they're also arriving quickly; sometimes faster than institutional policies and professional ethics frameworks can keep up with. Organisations are being asked to innovate while still safeguarding trust, care and responsibility.

One of the most intriguing (and for some, unsettling) applications of these technologies is what we call **AI afterlives**: using AI to 'revive' or simulate voices, faces or objects from the past. Sounds like something out of science fiction, right? But it's already happening. From deepfake historical figures giving museum tours, to chatbots and AI-curated relics, cultural heritage is stepping into a world where the past doesn't just sit in archives - it talks back.

At this stage, you might like to look at some examples via the link or QR code:



<https://www.syntheticpasts.com/workbookresources>



Maybe creating an AI afterlife is something you are already thinking about doing in your organisation, but you have questions. BIG ones:

- Are AI afterlives right for my institution?
- Does creating them align with our values?
- What opportunities and challenges do AI afterlives present?
- Where do we even start?

These are all great questions. Fear not, you are exactly where you need to be.

Introducing the Responsible AI Afterlives Workbook

This workbook was created as a practical companion: a space to pause, reflect and ask questions before jumping in. It was produced in consultation with cultural professionals, and is informed by a deep understanding of heritage work and responsibilities.

Who is it for?

This resource is designed for:

- Museums, heritage sites, galleries, libraries and archives who are, or might be, considering AI afterlife projects.
- Curators, producers, educators and digital teams exploring new tools.
- Technology companies and developers who want to better understand the values and priorities of heritage institutions.
- Community partners and stakeholders curious about how AI afterlives might shape their stories and heritage.

Don't worry, No technical expertise is required. The activities are self-led and adaptable for individuals, small teams or whole organisations.

What will it do?

The workbook has been designed to help you navigate this emerging landscape with care, curiosity and responsibility. It doesn't have all the answers (sorry!), but it will help you ask the right questions to explore whether AI afterlives are the right fit for your work.

It will help you to:

- Assess your needs – decide whether creating AI afterlives will be genuinely useful or if another approach works better.
- Think ethically – make sure your use of AI afterlives stays aligned with cultural, community and institutional values.
- Foreground responsible AI innovation – understand the issues raised by AI in more general terms, and approach it with care.
- Get started – with prompts, exercises and real-world examples.


Using the workbook will help you and your organisation build confidence to make informed, thoughtful decisions – ensuring AI enhances (rather than compromises) the integrity of your work.

How should you use it?

- Treat it like a workbook – write in it, circle things, sketch ideas in the margins.
- Use it solo or with colleagues – it works for reflection, team discussions, or workshops.
- Dip in where needed – you don't have to move through it in order; go to the sections that matter most right now.
- Expect prompts, not prescriptions – the exercises are here to spark thinking, not to give you yes/no answers.

This isn't a rulebook, it's a conversation starter. The aim is to balance curiosity with care, to give you confidence to say yes, no, or not yet to the creation of AI afterlives, and to approach them in ways that reflect your mission, your values and the communities you serve.

Are you ready? Let's get started.



Part 1: Values first

Bringing historical figures back to life with AI can be thrilling - but before you start summoning digital ghosts, it's worth asking: Should you?

As you explore that question, you'll naturally bump into the values and principles that steer your work, and those will help you see whether AI afterlives really belong in the picture.

Why does this matter? Because values act as your filter. Without values, it's easy to get swept up in novelty or pressure to innovate. With them, you have a clear way to check if a project truly aligns with your mission, your community and your responsibilities.


By capturing your professional values here, you create a touchstone you can keep coming back to. When the temptation is to chase novelty or when a project feels uncertain, this list will help you ask:

Does this align with what matters most to me professionally?

If you work for an institution or are thinking about using AI afterlives in an institutional setting, this work will likely be guided by an institutional mission, vision or set of values. Identifying these early on can help you answer:

Does this project reflect who we are as an institution, what we stand for, and the communities we serve?

You may also have an institutional AI policy which you can use to guide and safeguard your approach. Such a policy might help clarify things like copyright and image rights, which datasets and systems you can work with, and any approvals or risk assessments that might be necessary to do this work.



Worksheet 1: Your values map

This resource takes a values-based approach to AI afterlives, as cultural professionals have told us that values underpin the work that they do. This is where you'll think about the values you - and your institution - hold.

STEPS

- 01** Take a moment to look at the following professional values that might be relevant in heritage contexts.



- 02** Make a note of the ones that feel most important to you. Feel free to add others that come to mind. There's no right answer, only the principles that make your work meaningful.

The values that guide my work include:

- 03** If you have an organisational mission statement, vision or core values, think about that for a moment and make some notes if you wish:

Mission / Vision / Core Values:

- 04** Does creating an AI afterlife align with the above values?

YES - It aligns because:

NO - It does not align because:

UNSURE - We need more discussion because:

Keep these values in mind - you'll return to them.

Part 2: Why AI afterlives

The best AI projects don't replace human expertise - they amplify it, freeing up time for deeper storytelling, creativity and engagement.

Being clear about your goals means you can be accountable, and accountability is hugely important when it comes to working with AI in any capacity.

Before jumping in, ask:


What can AI do here that truly enhances our work?

Maybe you are hoping an AI afterlife project will uncover hidden connections in collections, trial a new approach to learning and engagement, create immersive experiences that 'bring history to life' or encourage more people through the door.

You should also ask yourself and colleagues:

What can an AI afterlife do that other forms of interpretation or engagement cannot?

This is a way of capturing the unique contribution of creating an AI afterlife; its USP (Unique Selling Point), if you like. Perhaps you are looking for it to support more dialogue and interaction, more personalised encounters with heritage or the generation of new material about the past.



Worksheet 2: Articulating goals

Every cultural institution has its own goals and objectives, and there may be a range of ways to achieve these. AI afterlives should add something that other approaches can't.

STEPS

01 Ask yourself (or discuss with your team):

What can an AI afterlife do for us?

- ☐ Provide personalised or adaptive interactions
- ☐ Reveal hidden connections
- ☐ Spark dialogue in new ways
- ☐ Create a sense of immediacy or presence
- ☐ Allow visitors to "converse" with history
- ☐ Animate voices, faces or memories in dynamic ways
- ☐ Blend fact and imagination to provoke reflection
- ☐ Reach audiences outside the physical space (online, remote, global)
- ☐ Make heritage feel contemporary and relevant
- ☐ Encourage empathy or deeper emotional interactions
- ☐ Bring more people through the door and/or increase engagement

Something else:

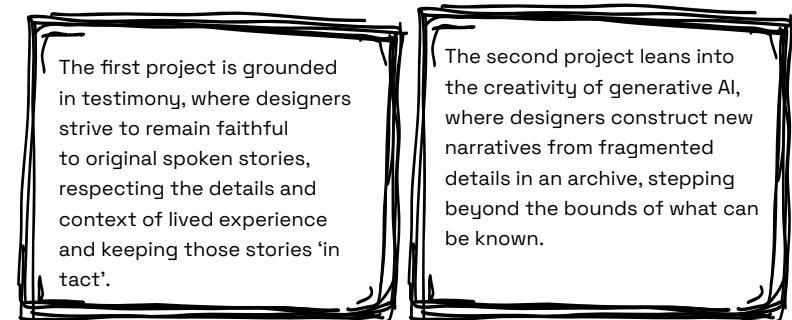
02 Can AI help you better achieve these outcomes than other approaches? Outline your thoughts below:

Part 3: Why not AI afterlives

Once you have spent some time thinking and chatting about your goals it is time to consider some of the risks involved.


In the creation of AI-generated afterlives the boundaries between reality, imagination and technology blur. AI offers powerful tools that invite playfulness, deception and hallucination - each acting as a catalyst for new forms of expression and exploration.

Take, for example, two (imaginary) contrasting AI afterlife projects:



Both projects encourage us to consider the role of illusion and storytelling in work with AI, and how they might coexist in a world where trust is hard-won and needs to be actively guarded.

If the balance tips the wrong way, you could risk reputational damage - never a good outcome, especially when thinking about investment and impact. This is why it's vital to be clear from the start about what success looks like for this work.



Worksheet 3: Weighing up the risks

Every new idea carries risks and AI afterlives are no exception. The goal isn't to avoid risk altogether, but to see it clearly, weigh it carefully and decide how to respond. By mapping risks openly, you strengthen transparency - one of the most important values when working with AI.

STEPS

- 01** Map out the risks you can foresee in relation to AI Afterlives.

- 02** Once you've mapped the risks, take a moment to think about what you can do about them:

- 03** Finish by weighing the big picture:

Do the goals identified in the previous section outweigh the risks you've identified?

Part 4: Ethics and responsible AI

Ethics isn't just a checkbox - it's the compass that keeps your AI afterlife journey on the right path. Considering ethics helps us slow down and make sure our choices are respectful, responsible and aligned with our role as custodians of heritage.

When working with AI afterlives in cultural heritage, some big questions to ask yourself are:

Who is shaping this technology?
Whose stories are made visible and whose are erased?
Who benefits, and who might be left out?

Consent matters. Remember that legal permission and ethical responsibility don't always align, and try to be mindful that custodianship is key.

Unlike the big tech companies who often shape these tools, cultural professionals are caretakers of memory, heritage and story, with a keen eye for:

Representation

When it comes to AI... AI can amplify biases just as easily as it can uncover hidden histories, shaping whose voices are heard.

Accuracy

When it comes to AI... Inaccuracies introduced by AI may distort how societies remember the past and imagine the future.


Access and Inclusion

When it comes to AI... If adopted uncritically AI risks creating new forms of exclusion that contradict institutions' mission of universal cultural access.

Sustainability

When it comes to AI... Data centres, software and generative AI tools consume energy. The environmental cost should be justified by the cultural value a project creates.

Take the long view. Sometimes the most ethical decision is to pause until your data is more robust, your partnerships are in place, or your institution is ready.



For more resources on responsible AI for cultural heritage, scan the QR code or head to www.syntheticpasts.com/workbookresources



Worksheet 4: Ethics in action

AI afterlives can open up exciting new ways of connecting with the past but they also raise tough questions. In considering the ethical implications of AI afterlives we want to highlight the values of responsibility, respect, dignity and trust.

STEPS

- 01** Explore the two scenarios that follow. They're not about right or wrong answers, they're about noticing the dilemmas, asking better questions and reflecting on what feels appropriate in your own context.

Scenario 1: Consent & voice

The family of a suffragette from the early 1900s has donated her diaries and these have been digitised. You're considering creating an avatar or chatbot that speaks "in her voice" based on these writings.

Guided questions:

- The family has given you the diaries, but do they fully understand how you want to use them? How is consent configured in this case?
- What kinds of questions and conversations do you anticipate the public would want to focus on?
- Does the chosen format impact how you feel about this scenario (e.g. whether the AI afterlife is a talking avatar/deepfake, or a text chatbot?) What might be lost (or gained) in users' interactions by choosing one format over another?

Potential ethical issues identified:

Scenario 2: Contested Objects

Your institution (which is based in the UK) holds an extensive collection of items connected to the Boer War. You'd like to create an online archive where people can "speak to" those objects, asking questions about their provenance and "experiences" of the war.

The collection includes uniforms, medals, photos and creative works, and a Boer bible taken from the home of a local family. The donors say you can use the items however you like.

Guided questions:

- On the basis of the collection, whose stories are likely to be told here and whose might be missing?
- What are the ethical implications of 'reviving' objects connected to colonial violence?
- Who benefits if this AI afterlife is created? Who might feel uncomfortable?
- Is consent from the donors of the items enough to proceed or are there others whose perspectives you should consider?

Think about who should be consulted or involved. This could include families, cultural groups, developers or expert advisors.

- 02** Reflecting on your own context, what ethical questions are likely to emerge? Who should you involve in the conversation?

Part 5: What to consider when designing responsible AI afterlives

You've explored values, weighed risks, tested ethics and begun to map the possibilities. In this section we introduce a series of considerations for working responsibly with AI afterlives.

Creativity & Care

Designing an AI afterlife isn't just about what the technology can do, it's about the kind of experience you want to create. Responsible design balances imagination with intention: it's playful enough to spark curiosity, yet thoughtful enough to respect the people, stories and communities involved.

You should also consider how the use of AI might reshape your creative processes and relationships within and beyond your institution, influencing who contributes, how work is made, and how meaning is shared with audiences and communities.

Community Impact

AI afterlives don't just affect your organisation, they ripple outward into the lives, memories and identities of the communities you serve. That's why being clear about who needs to be involved is essential. From your immediate team to communities or advisors, from funders to future audiences, each has a stake in how stories are told and retold.

Cultural institutions have worked hard in recent years to broaden narratives and challenge exclusions. Work with AI should build on that progress, not undermine it.

Accessibility and Inclusivity

An AI afterlife has the potential to open doors, making cultural heritage more accessible to people. But this needs to be done thoughtfully. Accessibility ensures that everyone, including users who are neurodivergent or who have disabilities, can engage with and benefit from your AI afterlife. Inclusive design not only meets ethical and legal responsibilities but improves usability and reach for all users.

Any interaction with an AI afterlife should feel simple and intuitive. The goal is to make it work for everyone it's intended for, not just those who are already comfortable with technology.

Data Integrity

An institution's values are also reflected in the way it handles data; their accuracy, reliability, maintenance and use. AI afterlife projects are grounded in data, and can only ever be as robust, inclusive, or transparent as those data that underpin them. Data integrity is a key consideration then and should inform decision-making across the life-cycle of the work.

Reflection and Iteration

Building in opportunities to reflect at varied points in your project is important. This will help to keep your project in alignment with your institution's values, but will also lead to more robust and imaginative work with AI afterlives.



Worksheet 5: Designing responsible AI afterlives

There's a lot to think about when it comes to creating AI afterlives. This is where we begin to move toward more practical steps.

STEPS

01

Consider the questions below and make notes to inform your next steps. If you have an AI afterlife project in mind, you can begin to think about how these questions might apply in that work. If you do not currently have a project planned, perhaps you can imagine one, or you can pick one of the examples in this workbook.

Creativity and Care

Who is this AI afterlife for, and how will it feel to those who interact with it?

What emotional impacts might arise (grief, distress, uncanny valley)?

What will users take away from an interaction (memories, insights or something more tangible)?

How will the boundaries of this experience be defined; what the AI afterlife will and will not do?

Can/should co-creativity with communities be integrated?

How will design decisions and trade-offs be documented?

Creativity and Care

Community Impact

What stories are being told? Who is telling them? Who might be silent/silenced?

Who benefits from this AI afterlife and who might feel uncomfortable about its creation?

Whose/which data are being used, and are permissions needed?

Which communities and stakeholders should be involved? How should their contribution be recognised and remunerated?

How might the intended and unintended legacies (the afterlives) of this work be understood? Could the AI afterlife be put to use in ways not yet anticipated?

Community Impact**Accessibility and Inclusivity**

Does this AI afterlife enhance access or create new obstacles? How can engagement with the AI afterlife be made more accessible in line with digital accessibility principles?

Where will people encounter it? (online perhaps or in an exhibition space). What challenges or opportunities does that context present?

What does inclusivity look like in this project? And who might need to be involved in order to make that happen?

Accessibility and Inclusivity**Data Integrity**

Which datasets and data partners will be used? What opportunities and challenges do those present? Are data legible, complete, and unbiased? If biases exist, how can they be acknowledged or addressed?

Does the data relate to objects or humans? Should these be approached differently?

How can transparency about data sources be structured into the experience?

What expertise is available internally, and what needs to be sourced externally?

What happens if new data comes to light after the AI afterlife has been created?

Data Integrity**Reflection and Iteration**

How can the effectiveness of this work be understood? How will 'success' be defined?

How will transparency and accountability be ensured?

How will user testing and consultation be built in? How will these insights impact the approach taken?

Do existing evaluation methods work for AI projects, or are new ones needed?

What is the plan for system failure or retirement?

Reflection and Iteration

Part 6: Your AI afterlives action plan

This workbook has not set out to give you a simple yes or no answer to the question of whether you should create AI afterlives. Instead, it has provided something more resilient; a structured way to pause, weigh options, and think through the implications of using AI to create afterlives in cultural heritage.

Over the previous four sections you have:

- Identified the professional and/or institutional values that guide your work.
- Considered the reasons for and against creating AI afterlives.
- Weighed the ethical challenges and responsibilities involved.
- Explored design considerations.

You should now have a clearer idea if AI afterlives are right for your institution, and if so what those AI afterlives might look like.

Now it is time to look forward. What will you take from this? What are the conversations and actions that need to happen next? If you choose to move forward, what is the first small step you will take?

Next steps can include:

- A commitment to discuss the questions raised here with your team,
- An action point to discuss any technology decisions with an appropriate person (e.g. in your IT, legal or ethics team),
- A reminder to remain realistic about long-term risks as well as short-term benefits,
- A prompt to obtain diverse expert input, including from relevant communities.

Success means reaching a decision that is clear, shared, and realistic; whether that is to create an AI afterlife, to pause, or to keep exploring.

Worksheet 6: Your AI afterlives action plan

Before you finish, take a moment to pause and capture where you stand right now. You do not need to make a final decision, but these prompts can help you see what feels right for your organisation at this point.

STEPS

- 01** Where do you stand?
Think about everything you have reflected on in this workbook. Which of these statements feels closest to your position?

We will not pursue an AI afterlife at this time.

- ☐ Next step: Capture reasons for this decision
- ☐ Action ideas: note conditions that might change the decision; file reflections for future reference

We are unsure and need to explore further.

- ☐ Next step: Identify the biggest open questions.
- ☐ Action ideas: Plan small experiments; set up discussions with IT/ethics/legal/community partners; or commission a scoping study.

We want to move forward with an AI afterlife project.

- ☐ Next step: Define clear goals and success criteria.
- ☐ Action ideas: Map out goals, risks, and mitigations; identify people to involve.

- 02** If you are planning to proceed with an AI afterlife project you should consider the following:

What skills do we currently have, and what might we need support with?

Who - in our organisation - needs to be part of this discussion?

- ☐ My immediate team
- ☐ IT / Digital specialists
- ☐ Legal / ethics advisors
- ☐ Community representatives
- ☐ Other:

Who - beyond our organisation - might we need to hear from?

Which resources or frameworks will guide us in this work? (You can find some helpful resources at www.syntheticpasts.com/workbookresources)

One goal and one risk

Keeping perspective

How will we balance creativity and care in this work?

- ☐ Schedule regular check-ins and reflection meetings.
- ☐ Capture, share and discuss successes and challenges as we go.
- ☐ Commit to transparency within - and beyond - the team about limitations and uncertainties.
- ☐ Plan for a regular cycle of user testing, evaluation and iteration.

Other plans:

Use worksheet 6 to guide your next conversations and decisions about AI afterlives. So whether you say yes, no, or not yet, your choice is both meaningful and responsible.

If you want to know more about the Synthetic Pasts project and the work we are doing on AI afterlives, you can visit www.syntheticpasts.com



Glossary of key terms

AI

Artificial Intelligence. Computer systems that learn from data to perform tasks like recognising images, answering questions or generating text.

Main types used in heritage:

- Machine learning: systems trained on data to recognise patterns (e.g. cataloguing).
- Natural language processing (NLP): systems that work with human language (e.g. chatbots).
- Generative AI: systems that create text, images, audio, or video (e.g. image or text generation based on existing dataset).
- Computer vision: AI that interprets images or video (e.g. identifying objects in digitised collections).

AI Afterlives

The use of AI to simulate or “revive” voices, faces, or objects from the past. In museums or archives, this might mean interactive historical figures, chatbots in an ancestor’s voice, or AI-generated archival material. (Scan the QR code for examples.)



Avatar

A digital representation of a person or character. They are used in museums, archives, or heritage sites to enable interaction with visitors. Avatars can appear as 3D characters, holograms, chatbots, or virtual guides and are often powered by AI or motion capture.

Chatbot

A program that uses AI to have conversations with users through text or speech. In museums, chatbots can answer visitor questions, guide tours, or simulate conversations with historical characters.

Dataset

A collection of data used to train or run an AI system. In heritage, datasets may be catalogues, digitised archives, image collections, or user-generated content. The quality, scope, and biases of a dataset directly shape the AI’s outputs.

Deepfake

AI-generated media (often video or audio) that imitates a real person’s likeness or voice.

Ethics

Principles guiding what is right, fair, and respectful in professional practice. For AI in heritage, ethics includes issues of consent, representation, bias, and care for communities and stories.

Generative AI


A branch of AI that creates new content (text, images, video, audio) based on patterns learned from large datasets. In museums, it can be used to generate exhibition text, images, or interactive reconstructions, but it may also “hallucinate” or invent.

Prompt Engineering

The skill of crafting clear and precise instructions (“prompts”) to get useful outputs from generative AI tools. For example: asking an AI to explain an object for a 10-year-old visitor vs. for an academic researcher.

Responsible AI

The approach of using AI in ways that are transparent, accountable, fair, and aligned with institutional values. In heritage, this means ensuring that projects enhance trust, respect communities, and do not distort history. Scan the QR code for more resources.
<https://www.syntheticpasts.com/workbookresources>



DOI: 10.5281/zenodo.17619346
CC BY 4.0 | © 2025 Synthetic Pasts | Designed by yello brick