



## Dialogue II

### **Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding**

By Evgeny Kalachikhin

**When AI can generate unlimited Holocaust imagery, protecting authenticity requires establishing methodological constraints rather than avoiding these technologies. Three principles - transparent marking of generated content, scholarly verification, and spatial grounding - could transform AI from a threat to authenticity into a tool for responsible visualisation. These constraints accept that some visual representation, even if imperfect, may serve memory better than leaving crucial narratives invisible to contemporary audiences - provided the imperfection is acknowledged and minimised through rigorous frameworks.**

## Introduction: Between Capability and Responsibility

Mykola Makhortykh and Maryna Sydorova's provocation identifies a critical challenge: how can we protect the authenticity of Holocaust evidence when facing a wave of AI-generated histories and memories? This question becomes particularly urgent when considering what remains hidden in archives. Vast amounts of Holocaust narratives exist only in textual form - academic research, survivor testimonies, institutional documents. These include stories of resistance, courage, and suffering that contemporary audiences may never encounter.

Younger generations increasingly engage with history [through visual media](#). For practitioners working with fragmentary visual evidence, the challenge is practical: should these narratives remain inaccessible? Or can methodological frameworks transform AI from a threat to authenticity into a tool for responsible visualisation? The answer to this question may also address another concern from the provocation: what could motivate the public to choose Holocaust-sensitive systems over unconstrained commercial platforms? Perhaps the very constraints that protect authenticity - transparency, verification, spatial grounding - could become the features that earn public trust.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

These questions emerge not from abstract concern but from urgent practical reality. As the provocation notes, AI-generated Holocaust content already floods digital environments - from [TikTok](#) videos depicting prisoner life in Auschwitz to [fake historical photographs](#) shared on social media platforms. Simultaneously, heritage institutions experiment with AI for preservation and education, from [Yad Vashem's named-entity recognition](#) to [interactive survivor testimonies](#). The distinction between these applications lies not in their technology but in their methodology.

Drawing primarily from visual practice - image generation and spatial visualisation - this response suggests that the North Star for AI in Holocaust memory might be found not in technological sophistication but in methodological constraint. The principles proposed may extend to other forms of AI-generated content, including text and audio, yet they emerge here from practical engagement with visualisation. Holocaust-sensitive AI systems could distinguish themselves through principled limitations that build trust - explicitly acknowledging provisional status, requiring expert validation, and anchoring generation in documented physical reality. The challenge facing the field is to demonstrate that these constraints enhance rather than hinder understanding, establishing frameworks that transform AI from a threat to authenticity into a tool for responsible memory work.

## The Documentation Gap and Its Consequences

Holocaust memory confronts an inherent challenge: the systematic [destruction of evidence](#) was integral to the genocide itself. The Nazis destroyed documentation, murdered witnesses, and dismantled sites of atrocity. What [remains is fragmentary](#) - testimonies from survivors, documents preserved by chance, photographs taken by perpetrators. For many historically significant sites and events, visual documentation ranges [from sparse to non-existent](#).

This absence creates a profound tension in [contemporary memory work](#). Younger generations increasingly engage with history through visual media, yet crucial narratives exist primarily in textual form - academic articles, survivor testimonies, archival documents. The provocation identifies this tension when discussing how AI is applied for "storytelling about the Holocaust", noting projects that transform testimonies into visual formats. However, the field has yet to establish clear principles distinguishing responsible visualisation from fabrication.

The consequences of this gap extend [beyond individual institutions](#). When heritage organisations cannot make their narratives visually accessible, they may cede ground to  
 © Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

actors operating without scholarly oversight or ethical frameworks. The provocation's examples - AI-generated images of Holocaust victims shared on social media, fake historical materials propagating antisemitic messages - demonstrate what can occur when visual generation operates without constraint. The question becomes not whether AI will be used to visualise Holocaust history, but who establishes the principles governing such use.

## Seeing Versus Not Seeing: A Practitioner's Perspective

From a practitioner's standpoint - particularly one coming from documentary filmmaking before engaging with AI - the ethical calculation may differ from purely preservationist approaches. Documentary and narrative cinema have long grappled with representing historical events that lack comprehensive visual documentation. Directors use re-enactments, dramatisations, and artistic interpretation to make history accessible. Well-known cinematic representations of the Holocaust faced [extensive criticism from scholars](#) regarding historical accuracy - specific details of environments, behaviours, and events. Yet their impact on public Holocaust awareness remains immeasurable.

This tension between scholarly precision and public engagement reflects a deeper question: what serves Holocaust memory more effectively - maintaining absolute fidelity to fragmentary evidence whilst accepting limited reach, or employing visualisation techniques that expand accessibility whilst acknowledging their provisional nature? The answer cannot be universal, but it should be intentional.

The argument here is not that accuracy matters less than impact. Rather, it recognises that all historical work involves degrees of interpretation and uncertainty. The question becomes whether AI-assisted visualisation, conducted under appropriate frameworks, might serve memory better than leaving crucial narratives inaccessible to visual learners. From a practitioner's perspective: seeing imperfectly may be preferable to not seeing at all - provided the imperfection is acknowledged and minimised through rigorous methodology.

This position consciously engages with longstanding debates about the limits of Holocaust representation. Critics from [Adorno](#) to [Lanzmann](#) have argued that certain aspects of the Holocaust should remain unvisualised - that representation risks aestheticisation, kitsch, or the domestication of horror into consumable imagery. These concerns deserve serious acknowledgement. Yet the risk of kitsch or inappropriate aestheticisation is not inherent to any technology, including AI. A film camera does not create exploitative imagery; a director does. Visual effects software does not produce tasteless spectacle; creative decisions do. AI, similarly, possesses no

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

agency - it generates nothing autonomously, makes no ethical choices, bears no responsibility. The human author bears full responsibility for every output, as with any act of creative or media production.

The concerns critics raise about AI "filling gaps" with inappropriate fabrications might be better understood as concerns about the rigour of human oversight rather than the nature of the technology itself. In professional production practice, every generated image results from deliberate choices: what to visualise, how to constrain generation, what to verify, what to acknowledge as provisional. From this perspective, the question of "seeing imperfectly versus not seeing at all" becomes less about technological capability and more about whether authors are prepared to accept responsibility for methodologically accountable visualisation.

This focus on authorial responsibility also clarifies the scope of the present argument. Video testimony already provides visually accessible Holocaust narratives for many learners. The challenge addressed here concerns a different gap: narratives preserved only in textual form - academic research, archival documents, secondary accounts - where no visual testimony exists. For these cases, the choice is not between AI visualisation and authentic video testimony, but between methodologically constrained visualisation and continued invisibility.

## Case Study: When Four Photographs Must Suffice

To ground these theoretical considerations in practical reality, consider a specific challenge: visualising a historically significant Holocaust-related site documented by only four low-quality archival photographs. The Pedanterie laundry in Bielsko-Biała, Poland, served as a clandestine contact point where Polish workers risked their lives maintaining communication between Auschwitz prisoners and their families. Detailed historical research by [Dr Jacek Proszyk](#), a [historian based in Bielsko-Biała](#), documents numerous acts of courage: messages concealed by a seven-year-old girl, food carefully placed on windowsills before prisoner transports, desperate moments of visual contact through windows, even a documented case where a prisoner became engaged to his fiancée in the laundry corridor on 18 March 1944.

These emotionally profound narratives exist almost exclusively in text. For contemporary audiences - particularly younger generations engaging primarily through visual media - these stories remain effectively invisible. The documentation gap creates a stark choice: accept this invisibility or develop methodologies for responsible visualisation.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

The “Pedanterie - the Auschwitz Laundry” project was developed to address this documentation gap through responsible AI-assisted visualisation. The project began with a fundamental question: could AI transform emotionally rich historical narratives into visual representations that capture their essence? Early experiments proved remarkably successful - working solely from Dr. Proszek's textual descriptions, AI generated images that local historians in Bielsko-Biała assessed as historically believable. These weren't mere illustrations but genuine transformations of text into visual memory.

This achievement was significant. For the first time, narratives trapped in academic texts became visible – families watching through windows, messages hidden by children, moments of human connection amid systematic dehumanisation. The fact that historians validated some of these purely text-based visualisations demonstrated AI's potential to bridge the gap between textual documentation and visual understanding. Yet the success rate, while meaningful, remained limited. Too many generated images, though emotionally compelling, lacked the grounding that would make them consistently reliable.

The project's evolution wasn't about abandoning text-to-image generation but enhancing it. The still-standing building offered an opportunity to increase the reliability of visualisations without sacrificing their emotional power. By documenting the actual architecture, the project could maintain what was already working - the transformation of narrative into image - whilst adding spatial constraints that would improve the percentage of historically grounded results. This experience revealed that responsible AI visualisation requires not just technology but systematic principles: making provisional status transparent, ensuring expert validation, and grounding generation in verifiable evidence where possible. These principles, emerging from practical necessity, suggest a broader framework for Holocaust-sensitive AI applications.

## Three Principles: Transparency, Verification, Constraint

Drawing from practical experience, three principles emerge as a potential foundation for Holocaust-sensitive AI applications. These principles do not claim to be exhaustive - the field remains nascent, and additional standards may emerge as practice develops. They represent necessary but not sufficient conditions: following them does not guarantee responsible visualisation but ignoring them substantially increases the risk of irresponsible outcomes. Each principle carries inherent limitations, and their application requires professional judgement rather than mechanical implementation. Practitioners must decide how prominently to disclose AI generation without undermining engagement, whose scholarly

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

assessment to trust when experts disagree, and what degree of spatial precision is required for different types of narratives. The principles outlined below offer a framework for such decisions, not a formula that eliminates the need for them.

### Transparency: Marking Provisional Status

AI-generated images or videos should ideally be clearly identified as such. This could extend beyond simple labelling to encompass methodological transparency: what sources informed the generation? What spatial or material constraints were applied? What aspects remain speculative?

This principle directly addresses the provocation's concern about authenticity and public trust. When TikTok users generate "[a day in Auschwitz](#)" content without identifying it as AI-produced, they may erode the distinction between historical documentation and creative interpretation. Heritage institutions using AI might maintain this distinction explicitly, treating generated content as provisional visualisation rather than discovered evidence.

The transparency principle faces practical challenges. In exhibition contexts, how prominently should AI generation be disclosed without undermining emotional engagement? In educational materials, how can we acknowledge provisional status whilst maintaining narrative coherence? These questions lack universal answers but suggest a need for institutional policies establishing clear standards.

Implementing this principle may also require distinguishing between transparency and explainability. [Recent research](#) suggests these represent different goals: transparency discloses that AI was used, whilst explainability renders the system's rationale and sources intelligible to audiences. A watermark or label fulfils the former but may fail the latter if the methodological process remains opaque. Responsible visualisation might aim beyond mere disclosure towards providing what could be termed "[scaffolding](#)" - accessible documentation of sources, constraints, and verification processes for those who seek it. Yet the limits of this approach should be acknowledged. Practitioners can ensure that such information remains accessible; they cannot guarantee audiences will engage with it. The challenge of cultivating critical engagement with AI-generated historical content extends beyond software design into the broader realm of digital literacy and Holocaust education - a challenge that institutions and educators must address alongside technological frameworks.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

## Verification: Scholarly Consultation

AI-generated historical content could benefit from expert review before public presentation. This verification serves multiple functions: identifying anachronisms, ensuring cultural sensitivity, confirming historical plausibility, and providing scholarly legitimacy.

For the Pedanterie project, collaboration with historians from the Institute of Urban Culture in Bielsko-Biała provided crucial verification. During fieldwork, extensive discussions with local historians assessed initial visualisations as "historically plausible" - a carefully calibrated judgement acknowledging both the limitations of available evidence and the responsibility to avoid fabrication. Historians provided detailed contextual analysis informing subsequent generations, establishing parameters for what elements could be interpretatively visualised versus what should remain strictly documentary.



**Figure 1:** Text-based AI visualisations recognised as believable by historians in Bielsko-Biała.

© Film University Babelsberg KONRAD WOLF

This principle directly addresses the provocation's question about Holocaust-sensitive AI systems: what distinguishes them from general tools? Scholarly verification could represent one clear distinction. Commercial AI services generate Holocaust-related imagery without expert consultation. Holocaust-sensitive systems could incorporate verification protocols - comparing generated content against authenticated historical records, identifying anachronisms through temporal databases, or requiring expert approval before public display. The distinction lies not in capability but in built-in accountability.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

The verification principle raises resource questions. Expert consultation requires time and funding. For smaller institutions or individual researchers, such resources may be limited. This suggests a potential role for coordinating bodies - perhaps organisations like the [International Holocaust Remembrance Alliance](#) - in establishing verification networks and standards.

### Constraint: Spatial and Material Grounding

Where possible, AI generation could be constrained by verifiable physical evidence. This might involve architectural documentation providing spatial boundaries, material analysis informing texture and appearance, or photographic evidence establishing visual parameters.

The Laundry project demonstrates this principle by using 3D documentation techniques that capture buildings from multiple angles to create precise digital models. These models then serve as spatial guides for AI generation - essentially teaching the system about the building's actual structure. The AI learns where walls exist, how spaces connect, and what physical relationships are possible, preventing it from generating historically impossible scenarios.

Beyond spatial boundaries, material properties provide another layer of constraint. AI systems can be trained to understand how different materials appeared in the 1940s - the texture of wool uniforms, the patina of aged metal, the quality of wartime fabrics. By studying both authentic photographs and carefully constructed reference materials, these systems develop an understanding of material properties that constrains their generation to historically plausible appearances.



**Figure 2:** Example of AI visualisation guided by 3D architectural documentation.  
© Film University Babelsberg KONRAD WOLF

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

These two approaches - spatial grounding and material understanding - work together to create complete historical scenes. The architectural framework defines where events could occur and how people could move through spaces, whilst material knowledge ensures that clothing, objects, and surfaces appear as they would have in the 1940s. A generated image might show a clandestine meeting in the laundry corridor, with the architecture determining the exact dimensions of the space and the placement of windows where families watched for prisoners. Meanwhile, the material understanding ensures authentic details: the rough texture of work clothes, the worn wooden floors, the quality of light filtering through industrial windows on a winter morning in occupied Poland.

Spatial and material grounding provides measurable constraints on AI generation, potentially transforming it from unbounded speculation to constrained interpretation. However, this principle faces limitations: not all Holocaust sites remain extant for documentation; not all materials survive for analysis; not all events occurred in spaces permitting spatial verification. The principle applies where possible but cannot universally govern all Holocaust visualisation.

Yet the absence of direct physical evidence does not necessarily render spatial grounding inapplicable. The provocation references "[Let Them Speak](#)", a project reconstructing experiences of victims who perished without leaving testimonies. Such cases might seem beyond the reach of spatial constraint. However, these victims lived in specific places - homes, neighbourhoods, synagogues, schools, workplaces - many of which survive or remain documented. Spatial grounding operates as a spectrum rather than a binary condition: from high-precision constraint where buildings remain intact, through moderate grounding using related documented locations, to broader contextual anchoring where only general spatial parameters can be established. The degree of uncertainty should be transparently acknowledged, connecting this principle directly to the first: where grounding becomes less precise, transparency about that imprecision becomes more essential.

## Confronting the "Distortion" Critique

The most serious criticism facing AI-generated Holocaust content centres on [historical distortion](#). Critics argue that generating imagery risks creating "false memories", undermining authentic testimony, and providing ammunition for deniers who claim existing evidence is similarly fabricated. These concerns demand serious engagement rather than dismissal.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

The response requires acknowledging several uncomfortable realities. Firstly, historical "accuracy" remains [elusive even in traditional scholarship](#). Historians construct narratives from incomplete evidence, making interpretative choices at every stage. Exhibitions selecting which photographs to display, textbooks choosing which events to emphasise, museums deciding how to present artifacts - all involve interpretative frameworks that shape historical understanding. This is not a weakness but an inherent aspect of historical work.

Secondly, the distinction between "authentic" and "generated" content proves less absolute than it appears - at least when AI generation operates within rigorous methodological frameworks. Digitised copies of historical documents - scanned, colour-corrected, enhanced - [already involve technological mediation](#). Transcriptions transform oral testimonies into text, losing vocal emotion and gesture; restored photographs require reconstruction of damaged portions; colourised historical images add interpretative layers absent from originals. Even [preservation itself](#) requires choices about what to preserve and how. These examples suggest we already accept various forms of mediation in historical representation.

Thirdly, the concern about false memories, whilst legitimate, must be balanced against actual public engagement with Holocaust memory. Research cited in the provocation indicates troublingly [low levels of Holocaust knowledge](#), particularly among younger generations. If scholarly commitment to absolute accuracy results in narratives remaining inaccessible to these audiences, we risk preserving precision whilst losing relevance. The question becomes not whether to accept imperfection, but how to manage it responsibly.

The counterargument is not that accuracy doesn't matter - it matters immensely. Rather, it's that the choice facing practitioners often isn't between accurate and inaccurate representation, but between imperfect visualisation and no visualisation. Consider again the Pedanterie laundry: without AI-assisted visualisation, these narratives of resistance remain confined to academic texts. With it - conducted under frameworks of transparency, verification, and constraint - they become accessible to broader audiences who might never engage with scholarly articles but will watch a museum installation or educational documentary.

This position accepts risk whilst arguing that managed risk serves memory better than absolute caution. Major cinematic Holocaust narratives, despite their scholarly critiques, profoundly [shaped public Holocaust consciousness](#). Their emotional impact, their accessibility, their ability to make distant history feel immediate - these qualities advanced memory work in ways that scholarly accuracy alone could not achieve. They represented

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

managed risk: not documentary footage, but historical drama employing artistic interpretation within carefully researched frameworks.

The same logic applies to AI-assisted visualisation. It represents managed risk: not authentic documentation, but provisional interpretation constrained by architectural evidence, verified by scholars, and transparently marked as such. The question becomes whether the field can establish robust enough frameworks to manage this risk responsibly - or whether attempting to do so opens doors better left closed.

## The Corridor of Memories: From Assets to Experience

The practical application of these principles could extend beyond generating individual images to creating immersive narrative experiences. The "Corridor of Memories" concept, developed as part of the Pedanterie project, demonstrates how architectural documentation and AI-generated content might combine into spatial storytelling.

This approach uses three-dimensional scanning techniques not merely to document architecture but to define the spatial structure of immersive installations. The captured model provides the "canvas" onto which historical narratives - combining authentic archival materials with verified AI-generated reconstructions - could be projected. Visitors move through space, experiencing layered visual narratives grounded in authentic architecture.



**Figure 3:** The Corridor of Memories installation concept.

© Film University Babelsberg KONRAD WOLF

This format could address several challenges identified in the provocation. It makes textual narratives visually accessible whilst maintaining spatial authenticity. It combines preserved evidence with interpretative reconstruction, potentially distinguishing between them

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

through presentation design. It creates emotionally engaging experiences whilst acknowledging the provisional nature of generated content.

Importantly, this approach may democratise advanced content creation. It requires no specialised infrastructure beyond photogrammetric equipment and consumer-grade computing resources. Cultural institutions lacking resources for traditional exhibition design might nonetheless create compelling narrative experiences. This accessibility matters particularly for sites outside major memorial institutions - places where local significance might not warrant massive investment but where stories deserve preservation and presentation.

## Competing With General-Purpose AI: The Challenge of Accessibility

The provocation raises a crucial question about practical adoption: what could motivate the public to rely on Holocaust-sensitive AI systems rather than general-purpose tools? After exploring methodologies for responsible visualisation, we must confront an uncomfortable reality: constrained, verified, institutionally-mediated AI systems may not match the accessibility of commercial platforms.

Mainstream generative AI will generate Holocaust-related content instantly, without verification, without spatial constraints, without transparent methodology. Users seeking visual interpretations of Holocaust narratives face a choice: wait for institutions to develop verified content under rigorous frameworks or generate it themselves using readily available tools. The latter option requires no expertise, no institutional affiliation, no waiting period.

Holocaust-sensitive systems compete at a potential disadvantage. By design, they might need to be slower, more constrained, more transparent about limitations. They cannot promise unlimited generation. They cannot provide instant results. They may require acknowledging uncertainty and provisional status.

The motivation for using such systems therefore might not derive from convenience. It could derive from trust - trust that verified, constrained, transparently-marked content serves understanding better than unrestricted generation.

It is worth clarifying what trust means in this context. Trust is not directed at a technology but at the author or institution employing it. Cinematic representations of the Holocaust

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

illustrate this point. Films like *Schindler's List* faced scholarly criticism for aestheticisation and historical liberties, yet earned public trust not because of the medium but because of directorial accountability - Spielberg's reputation, consultation with historians, acknowledgement of the work's interpretative nature. The same camera equipment in different hands would not command the same trust. AI operates similarly: a heritage institution with scholarly oversight and transparent methodology may earn trust that an anonymous content creator using identical tools cannot. The technology remains constant; the accountability differs. This suggests that Holocaust-sensitive AI systems build trust not through superior algorithms but through demonstrated authorial responsibility - the same foundation that has legitimised cinematic, literary, and artistic Holocaust representation for decades.

Building such trust may require demonstrating that methodological rigour produces qualitatively different outcomes: visualisations that connect to authentic spaces, interpretations validated by scholarly expertise, content that acknowledges rather than obscures its provisional nature. This suggests several practical implications. Holocaust-sensitive AI systems might make their verification processes visible, showing rather than merely claiming scholarly involvement. They could provide rich contextual information explaining how visualisations were constrained and what remains uncertain. They might offer clear comparisons between verified institutional content and unconstrained commercial generation, helping users understand why constraints matter.

Furthermore, heritage institutions might actively engage with commercial platforms rather than treating them as competitors to ignore. This could involve developing educational materials explaining the risks of unconstrained generation, creating toolkits for educators wanting to use AI responsibly, or establishing certification programmes for Holocaust-related AI content meeting verification standards.

The competition with commercial platforms may not be won through superior technology - general-purpose models will always offer broader capabilities. It might be won through demonstrating superior methodology, building trust through transparency, and helping audiences understand why constraints could serve memory better than unlimited generation.

## Towards Regulatory Frameworks

The provocation also asks whether AI use in Holocaust memory should be regulated. From a practitioner's perspective, some form of regulation - or at minimum, widely-adopted

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

standards - may be inevitable and necessary. The question is what form such regulation might take and who should establish it.

Heavy-handed regulatory approaches risk stifling legitimate innovation whilst proving difficult to enforce across jurisdictions. More promising might be industry-led standard-setting, similar to [museum accreditation programmes](#) or archival [best practices](#). Organisations like the [International Holocaust Remembrance Alliance](#), working with heritage institutions, technology developers, and scholars, could establish certification programmes for Holocaust-related AI applications.

Such standards might address several key areas. First, content marking requirements: establishing universal standards for identifying AI-generated historical content across platforms and contexts. Second, verification protocols: defining what constitutes adequate scholarly review for different types of content. Third, transparency requirements: specifying what methodological information should accompany generated content. Fourth, spatial grounding standards: establishing when architectural or material verification should be required.

Certification programmes could function similarly to ethical review boards for human subjects research. Institutions or individuals planning Holocaust-related AI projects would submit proposals detailing their methodology, verification plans, and transparency practices. Certified projects could display standardised marks indicating compliance with established principles. This approach might balance innovation with oversight, providing frameworks without imposing rigid restrictions.

However, such regulatory frameworks risk creating unintended consequences. Overcautious standards could lead to censorship of legitimate Holocaust education and commemoration efforts. The tendency toward risk-averse overcorrection - already visible in how commercial AI platforms [handle sensitive topics](#) - might extend to Holocaust memory, where blanket restrictions block responsible academic research, survivor testimony projects, and artistic memorial works alongside actual problematic content. When platforms and institutions prioritise avoiding controversy over enabling meaningful engagement, they often implement self-censorship that extends far beyond ethical requirements. The challenge lies not in establishing standards but in calibrating them: too loose, and they fail to prevent distortion; too strict, and they stifle the very memory work they aim to protect.

This limitation suggests that regulation alone may not solve the authentication crisis. It might need to combine with education - helping audiences critically evaluate AI-generated

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

historical content, understand verification indicators, and recognise the difference between constrained institutional generation and unconstrained commercial use.

## The Economic and Institutional Challenge

The provocation notes that Holocaust institutions often lack capacity for pushing technological innovation in AI. This constraint appears both financial and conceptual. Financially, verification processes, expert consultation, and methodological development require resources that preservation-focused institutions may not possess. Conceptually, shifting from preservation to content creation represents a fundamental identity transformation that many institutions may resist.

However, the economic argument could be inverted. AI-assisted content creation, conducted under appropriate frameworks, might generate new revenue streams that justify digitisation investments. When documentation of any kind serves only archival purposes, it represents pure cost. When the same documentation enables production-ready content for documentaries, educational materials, museum installations, and virtual experiences, it could become an economic asset.

This transformation might require institutional mindset shifts. Museums and archives could reconceive themselves not merely as preservation custodians but as content providers serving multiple markets: education, tourism, media production, interactive experiences. AI technologies, properly implemented, may facilitate this transformation by dramatically reducing the cost and complexity of generating production-ready content from archival materials.

The Pedanterie project, for example, attempts to explore this potential. Photogrammetric documentation serves immediate preservation needs whilst simultaneously enabling spatial installations, educational content, potential documentary use, and virtual experiences. A single capture investment generates multiple outputs, each potentially sustaining itself economically whilst advancing memory work.

However, realising this potential requires overcoming institutional conservatism. Many cultural organisations remain anchored to traditional preservation paradigms, viewing digitisation as an endpoint rather than beginning. Transformation requires demonstrating successful models, providing practical frameworks, and helping institutions understand how content creation can complement rather than compromise preservation missions.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

## Limitations and Uncertainties

Yet there remain unresolved problems of an almost philosophical character. First and foremost, AI-generated content, even under rigorous frameworks, cannot achieve absolute historical accuracy. Generative models occasionally produce "hallucinations" - plausible-seeming details without historical foundation. Verification processes may catch many such issues but cannot guarantee perfection. Spatial grounding constrains but does not eliminate interpretative speculation.

The three principles proposed in this response also face specific practical limitations. Transparency through labelling represents a necessary minimum, yet for highly affective content - imagery designed to evoke strong emotional responses - disclosure alone may not prevent emotional manipulation. The author must judge not only how to label content but whether certain affective approaches are appropriate at all. Verification encounters difficulty when scholars disagree; the author must decide whether to proceed only where consensus exists or to acknowledge disputed interpretations transparently. Spatial grounding varies in relevance depending on the narrative: some visualisations demand architectural precision, whilst others - depicting internal emotional states or abstract experiences - may require different forms of constraint entirely. In each case, the principle provides a framework, but professional judgement determines its application.

Furthermore, the emotional impact of visualisation carries risks beyond factual accuracy. Making Holocaust narratives more accessible and emotionally engaging might inadvertently trivialise suffering, reduce complex history to simple narratives, or provide aesthetic pleasure from representations of atrocity. These concerns deserve serious consideration rather than dismissal.

The field also lacks long-term studies of how AI-generated historical content affects learning, empathy, and historical understanding. Do visualisations enhance engagement or create false confidence in understanding? Do they complement authentic materials or replace them in public consciousness? Do they strengthen historical memory or contribute to its dilution? These questions await empirical investigation.

Additionally, the sustainability of verification frameworks remains uncertain. As AI technology evolves rapidly, maintaining scholarly expertise adequate for reviewing generated content becomes increasingly challenging. The resource requirements for verification may prove prohibitive for widespread adoption. Standards developed today may prove inadequate for tomorrow's capabilities.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

These limitations suggest not abandoning AI-assisted visualisation but proceeding with humility, maintaining transparency about uncertainty, and continuously evaluating impact rather than assuming positive outcomes.

## Conclusion: Defining the North Star

Returning to the provocation's central question: the North Star for AI in Holocaust memory might lie in accepting constraints that general-purpose models refuse. Holocaust-sensitive AI systems could distinguish themselves not through superior technology but through methodological restraint - transparent marking of generated content, scholarly verification, spatial grounding where possible, and explicit acknowledgement of provisional status.

This North Star suggests navigating between competing imperatives: making history accessible whilst maintaining accuracy, employing new technologies whilst respecting traditional scholarship, generating content whilst acknowledging uncertainty. It suggests accepting that seeing imperfectly might serve memory better than not seeing at all - provided the imperfection is acknowledged and minimised through rigorous methodology.

Practically, this vision could require several developments. Holocaust museums and archives might evolve from preservation-focused custodians to active content creators, recognising that digitisation investments can generate both archival and creative returns. International remembrance organisations might establish standards balancing innovation with oversight. Technology developers could build verification and transparency into system architecture rather than treating them as optional features. Holocaust educators might help audiences understand why constraints matter, building critical literacy about AI-generated memory content.

The competition with unconstrained commercial platforms may not be won through prohibition or superior technology. It might be won through building trust - demonstrating that verified, constrained, transparently-marked content serves understanding better than unlimited generation. This could require making verification processes visible, providing rich contextual information, and actively engaging rather than avoiding commercial platforms.

The wave of AI-generated Holocaust content identified in the provocation appears inevitable. The field faces a choice: establish robust frameworks for responsible visualisation now or cede the field to unregulated generation that will proceed regardless. The North Star might guide not towards avoiding AI but towards constraining it through principles that distinguish responsible practice from fabrication.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>

From a practitioner's perspective, the imperative remains: remembering may require seeing. When crucial narratives exist primarily in textual form, inaccessible to contemporary audiences, responsible visualisation might serve memory better than cautious invisibility. The question is not whether to use AI for Holocaust memory, but whether the field can establish and maintain frameworks that transform technological capability into responsible practice - frameworks that future generations will judge by whether they advanced understanding or contributed to its erosion.

## Questions for Discussion

1. Should heritage institutions establish formal certification programmes for Holocaust-related AI applications, similar to museum accreditation systems? Who should have authority to define certification standards?
2. How can Holocaust-sensitive AI systems compete with the accessibility of general-purpose tools when, by design, they must be slower and more constrained? What motivates users to accept these limitations?
3. When spatial documentation of Holocaust sites is impossible (sites destroyed, access restricted, safety concerns), what alternative grounding mechanisms can constrain AI generation whilst maintaining transparency about limitations?
4. How should institutions balance the goal of making Holocaust narratives visually accessible with risks that visualisation might trivialise suffering or provide aesthetic pleasure from representations of atrocity?
5. What empirical research is needed to understand how AI-generated historical content affects learning, empathy, and historical understanding? How can we measure impact beyond immediate engagement metrics?

---

**Evgeny Kalachikhin** is an independent Film Director and XR Artist, as well as Art Director and Academic Researcher at the Creative Exchange Studio, Film University Babelsberg KONRAD WOLF, Germany, where he focuses on AI-assisted cultural heritage visualisation and immersive media technologies. He is co-author of the “Pedanterie - the Auschwitz Laundry” project.

© Evgeny Kalachikhin (2026)

Kalachikhin, Evgeny (2026) Dialogues II – Response I: Holocaust-Sensitive AI Should Earn Trust Through Accepting Three Commitments: Transparency About Sources, Scholarly Verification, and Spatial Grounding. *Digital Memory Dialogues* 2(2).

<https://dx.doi.org/10.20919/ITWS2142>