Generative AI as an Adaptive Narrative Layer for Cultural Heritage: Towards Responsible Personalized Storytelling

Liliana Ardissono[1,†], Angelo Geninatti Cossatin[1,†], Noemi Mauro[1,†] and Claudia Picardi[1,†]

[1]University of Turin, Corso Svizzera 185, Torino, 10149, Italy

Abstract

Recent advances in Generative Artificial Intelligence have introduced unprecedented possibilities for dynamically producing personalized content. In cultural heritage contexts, these technologies enable a shift from static interpretive materials toward adaptive, context-sensitive narrative experiences. Generative systems promise increased engagement and accessibility for museums, archaeological sites, historic environments, natural heritage locations, and digital archives. However, their use raises significant epistemic and ethical concerns. Cultural heritage interpretation is not merely informational, but deeply connected to identity, memory, and institutional authority. This paper proposes a vision of Generative AI as an adaptive narrative layer embedded within curated heritage systems. Rather than replacing curatorial expertise, generative systems should operate within bounded, knowledge-grounded frameworks that preserve authenticity and trust. We outline key research challenges related to user modeling, contextual adaptation, epistemic integrity, governance, and evaluation, and propose some research challenges and directions for responsible generative personalization in cultural heritage.

Keywords

Cultural Heritage, Storytelling, Generative AI, Personalization

1. Introduction

Cultural heritage interpretation has historically relied on static media such as exhibit labels, printed guides, audio tours, and fixed digital descriptions. Large-scale digital platforms such as Europeana [1] have made heritage collections largely accessible, but their interpretive content remains predominantly static and uniform. Although digital technologies have introduced interactive displays and mobile applications, most heritage content remains pre-authored and uniform, offering limited responsiveness to individual visitors' interests, backgrounds, or situational contexts. Personalization research has begun to address this limitation through recommender systems and adaptive navigation interfaces [2, 3], possibly integrating user-generated content to enrich cultural heritage information [4]. Some of these systems employ location-based interactive narratives that use gamification to engage users; see, e.g. [5, 6, 7]. However, with few exceptions, such as [8, 9], they operate by selecting predefined content rather than dynamically constructing interpretive narratives that consider simultaneously the cultural and geographical context surrounding the user and the user's interests and goals.

The emergence of Large Language Models (LLMs) and related generative technologies introduces a qualitative shift. Instead of retrieving and presenting fixed texts, systems can now generate narratives in real time, adapting explanations, thematic emphasis, and rhetorical style to the user and context [10, 11, 12, 13]. In cultural heritage settings, such capabilities suggest the possibility of applying interactive digital storytelling across diverse environments, including museums, archaeological sites, historic cities, natural parks, virtual reconstructions, and digital archives. For example, Ferracani et al. [14] generate personalized fictional narratives set in real heritage locations, combining GPT-4 with AI-generated illustrations to increase visitor engagement.

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden

[†]These authors contributed equally.

liliana.ardissono@unito.it (L. Ardissono); angelo.geninatticossatin@unito.it (A. Geninatti Cossatin); noemi.mauro@unito.it (N. Mauro); claudia.picardi@unito.it (C. Picardi)

0000-0002-1339-4243 (L. Ardissono); 0009-0007-5378-7061 (A. Geninatti Cossatin); 0000-0001-8234-3266 (N. Mauro); 0000-0003-1170-1913 (C. Picardi)

Yet, cultural heritage is a domain in which interpretive authority, authenticity, and historical accuracy are central. Thus, the integration of Generative AI (which is subject to hallucinations) into heritage experiences requires careful consideration.

We argue that Generative AI should be conceptualized as an adaptive narrative layer embedded within curated knowledge infrastructures, rather than an autonomous storyteller. We articulate principles for responsible generative personalization and outline a set of research challenges relevant to the PATCH community.

In the following, Section 2 discusses the transition from static content to adaptive narrative generation, and Section 3 outlines some related challenges. Sections 4 and 5 discuss how such challenges could be faced and propose an agenda for that. Section 6 concludes the paper.

2. Transitioning from Static Content to Adaptive Narrative Generation

Digital personalization in cultural heritage has traditionally focused on recommending artifacts, tailoring tour paths, or filtering collections according to user preferences. See [15, 2, 3] for some surveys on these topics. Narrative content itself has remained largely static. Even when delivered through mobile applications, haptic user interfaces, or interactive kiosks, interpretive materials are typically drawn from a predefined repository of texts created by curators or educators [16, 17].

Generative AI enables a deeper form of adaptation at the level of narrative construction [18]. Moreover, multimodal generative models can integrate visual, textual, and spatial information into immersive heritage experiences [19, 20, 14]. Thanks to Generative AI, the presentation of information can be adjusted to different levels of expertise, reoriented toward specific thematic interests, or framed in alternative narrative styles with substantially reduced authoring effort. For example, a visitor interested in architectural techniques might receive a materially focused account of a building, while another interested in social history might encounter a narrative centered on daily life. A child might receive a simplified child-oriented story, while a domain expert might receive a technical presentation. These variations can be generated dynamically rather than pre-authored in multiple versions or implemented with less flexible match-making approaches as happened in the past [21, 22, 23]. Moreover, generative systems can be used to engage in conversations with users to answer specific queries.

This shift suggests reconceptualizing Generative AI as a mediation layer between curated heritage knowledge and individual users. Such a layer would draw upon verified databases, ontologies, and institutional archives while adapting narrative output in response to user profiles, interaction histories, and contextual factors such as location or group composition. Importantly, this approach does not eliminate curatorial authority; instead, it introduces a flexible interface for accessing curated content.

3. Challenges of Using Generative AI in Personalized Storytelling

Generative AI can increase engagement by constructing narratives aligned with individual interests and interaction patterns. Rather than delivering fixed descriptions, AI-driven storytelling systems can emphasize themes, characters, or historical dimensions that resonate with a particular visitor, fostering stronger emotional connections and sustained attention.

However, generative AI introduces risks in cultural heritage contexts. Large language models are known to produce hallucinations, i.e., plausible but inaccurate statements generated without grounding in verified sources. In heritage interpretation, such inaccuracies can undermine institutional credibility and distort public understanding of history. For example, in a study focused on the application of ChatGPT-4 in the cultural heritage domain, Spennemann shows that it has limited capability to synthesize and present complex concepts [24]. More generally, generative AI cannot guarantee, by itself, the quality and accountability of the information provided because it can extract data from unreliable sources [25, 26, 27, 24]. Retrieval-Augmented Generation has been proposed as a general strategy to anchor language model outputs in verified sources and reduce hallucinations [28]. Its application to cultural heritage contexts, discussed in Section 4, remains an active area of investigation.

Personalization itself also raises epistemic concerns. By tailoring narratives to user interests, systems may unintentionally overemphasize certain aspects of heritage while marginalizing others, creating personalized filter bubbles that have been frequently observed in the literature on recommender systems [29]. Moreover, complex historical events might be simplified, controversial perspectives might be filtered, and interpretive tensions might be reduced for the sake of coherence or engagement. Over-personalization risks transforming shared cultural narratives into individualized and potentially distorted accounts.

Bias represents another critical concern [30]. Generative systems trained on large corpora may reproduce dominant cultural perspectives while neglecting marginalized voices. In contexts involving colonial histories, contested heritage, or traumatic memory, biased or insensitive narrative generation can cause harm.

Finally, the introduction of autonomous generative systems can be perceived as weakening curatorial authority. Cultural heritage institutions operate as stewards of collective memory. Thus, any technological mediation must preserve institutional governance and interpretive accountability.

4. Towards Responsible Generative Personalization for Storytelling

To address these challenges, we propose that Generative AI in cultural heritage should operate according to principles of bounded generation, knowledge grounding, and curator-in-the-loop governance.

Bounded generation requires constraining narrative output through structured access to verified knowledge sources. Retrieval-Augmented Generation techniques can anchor responses in curated databases and ontologies. Systems should maintain traceability between generated text and underlying sources, enabling verification and transparency [31, 32].

Knowledge grounding must be complemented by explicit mechanisms for uncertainty representation and factual consistency checking. Generated narratives should distinguish between established historical facts, scholarly interpretations, and speculative reconstructions. Where uncertainty exists, systems should communicate it clearly rather than presenting speculative content as authoritative, as failure to do so may foster overreliance on generated narratives [33].

Curator-in-the-loop governance ensures that institutions retain interpretive control. Curators should define narrative boundaries, approve knowledge sources, configure personalization parameters, and monitor system outputs [17]. Rather than replacing human expertise, Generative AI should augment curatorial practice by providing flexible mediation within predefined constraints [34].

5. Research Challenges and Agenda

Positioning generative AI as an adaptive narrative layer within cultural heritage raises research challenges that extend beyond traditional personalization and recommender systems. Instead of simply selecting relevant content, generative systems dynamically construct interpretive narratives that must remain coherent, trustworthy, and ethically responsible. The research dimensions summarized in Table 1 outline key directions and possible research questions for investigating this emerging paradigm.

Narrative Modeling. Generative systems must move beyond producing isolated artifact descriptions toward constructing coherent narrative trajectories spanning multiple objects, locations, or moments of a visit. This requires integrating narrative planning with personalization mechanisms so that adaptive explanations remain embedded within a consistent interpretive trajectory.

Interpretive Framing and Perspective. Cultural heritage narratives can be interpreted through different lenses, such as social, political, technological, or environmental perspectives. Generative systems capable of modulating these interpretive frames can provide richer and more relevant storytelling experiences for diverse audiences. However, enabling flexible reframing without distorting historical integrity remains an open challenge. This is delicate as Generative AI tends to be sycophant [35].

Knowledge Grounding and Traceability. Because cultural heritage institutions are expected to provide reliable and verifiable information, generative storytelling systems must remain tightly anchored

Research Dimension	Key Research Questions
Narrative Modeling	How can generative systems balance adaptability with global narrative structure across artifacts and sites?
Interpretive Framing & Perspective	How can generative systems modulate interpretive lenses (e.g., social, political, technological, environmental) without distorting historical integrity?
Knowledge Grounding & Traceability	How can generated narratives remain verifiable and anchored in curated sources? How can systems expose evidentiary links transparently?
Historiographical Complexity & Uncertainty	How can generative storytelling communicate uncertainty, debate, and incomplete evidence responsibly?
Bias, Inclusion & Cultural Sensitivity	How can generative storytelling avoid reproducing dominant narratives and marginalizing alternative voices?
Educational & Cognitive Impact	Does adaptive generative storytelling improve understanding, retention, or critical thinking compared to static interpretation?
Transparency & User Awareness	How should systems communicate that narratives are AI-generated and dynamically constructed?
Responsible Dialogue Management	How can conversational generative systems dynamically adapt to user inputs while preventing the propagation of misinformation, bias, or unsafe interpretations in cultural heritage narratives?
Ethical Limits of Personalization	When should personalization be constrained in generative storytelling (e.g., traumatic memory, contested heritage)?
AI–Curator Collaboration and Governance.	How can interactive tools and workflows be designed to enable curators to efficiently supervise, audit, and steer generative storytelling systems without excessive cognitive or operational burden?

Table 1
Research questions for generative AI in personalized cultural heritage storytelling.

to curated knowledge sources such as archives, scholarly databases, and institutional collections. Retrieval-Augmented Generation and provenance-aware architectures represent promising approaches, but further research is required to design mechanisms that transparently link generated narratives to their evidentiary sources.

Uncertainty and Historiographical Complexity. Historical interpretation often involves incomplete evidence, scholarly debate, and evolving perspectives. Generative storytelling systems should therefore be able to communicate uncertainty, disagreement, or interpretive plurality without collapsing them into overly simplified or authoritative narratives.

Bias, Inclusion, and Cultural Sensitivity. Large Language Models might reproduce dominant cultural narratives while marginalizing alternative voices or underrepresented communities [36]. In heritage contexts that involve colonial histories, contested identities, or traumatic memories, such biases may have significant social implications. Research is needed on methods for systematically auditing heritage-specific corpora and generated outputs for representational gaps, and on tools that support curators in identifying and addressing such biases within generative storytelling pipelines.

Educational and Cognitive Impact. Adaptive storytelling is often assumed to improve visitor engagement. However, its effects on learning outcomes and interpretive depth remain largely unexplored. Comparative studies between static interpretation and generative storytelling interfaces are needed to assess potential gains in knowledge retention, understanding, and engagement.

Transparency and User Awareness. Visitors should be able to recognize when narratives are dynamically generated and understand the role of AI in shaping the interpretive experience. Designing interfaces that communicate algorithmic mediation and the assumptions that the system has used to

personalize the content, without disrupting narrative immersion, remains a challenge. Transparency mechanisms should promote informed engagement while supporting appropriate calibration of trust.

Responsible Dialogue Management. Unlike static recommendation settings, conversational interaction introduces unpredictable and evolving user inputs that may steer narrative generation toward inaccurate, biased, unsafe, or epistemically weak outputs. Research is needed on dialogue management architectures that preserve responsible system behaviour under user influence, combining robust intent understanding, bounded generation, source-grounded responses, and uncertainty handling.

Ethical Limits of Personalization. Although personalization can increase engagement, excessive adaptation risks filtering out challenging perspectives or oversimplifying complex historical realities. In contexts involving sensitive heritage topics, such as conflict or trauma, personalization may need to be constrained to ensure respectful and responsible interpretation.

AI–Curator Collaboration and Governance. While curators should retain interpretive authority, source control, and responsibility for narrative boundaries, embedding human expertise into adaptive AI pipelines can be operationally complex and resource-intensive. We need scalable interaction models that enable curators to supervise outputs, define editorial constraints, inspect provenance, correct system behaviour, and iteratively refine personalization strategies without excessive workload, as investigated in [34].

6. Conclusion

Generative AI introduces the possibility of moving cultural heritage interpretation beyond static descriptions toward dynamically generated *personalized storytelling*. Instead of presenting predefined explanations, generative systems can construct narratives that adapt to visitors' interests, knowledge levels, and situational contexts, potentially supporting richer and more engaging encounters with cultural heritage across museums, archaeological sites, historic environments, and digital collections.

However, storytelling in cultural heritage is not merely a matter of information delivery. Historical narratives contribute to the construction of collective memory and cultural identity, and therefore require careful attention to authenticity, interpretive responsibility, and institutional authority. The use of generative AI in this domain must therefore balance personalization and narrative flexibility with mechanisms that preserve historical integrity and public trust.

In this paper, we proposed conceptualizing generative AI as an *adaptive narrative layer* that enables personalized storytelling while remaining grounded in curated heritage knowledge. Within this perspective, generative systems can mediate between institutional sources and individual visitors, dynamically generating narratives while operating within bounded, knowledge-grounded frameworks. This approach emphasizes the role of generative AI as a complement to curatorial expertise.

Building on this vision, we identified relevant research challenges for the application of generative AI in personalized cultural heritage storytelling, highlighting challenges related to narrative modeling, interpretive framing, knowledge traceability, representation of uncertainty, bias and inclusivity, educational impact, transparency, and the ethical limits of personalization. Addressing these challenges requires interdisciplinary collaboration across recommender systems, natural language generation, human-computer interaction, digital humanities, and museum studies.

By carefully integrating generative technologies with curatorial knowledge and responsible design principles, generative AI has the potential to support more engaging, inclusive, and critically informed forms of personalized storytelling in cultural heritage.

Acknowledgments

Our University has partially funded this work through the project "Grant for Internationalization".

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT[1], Claude[2] and Grammarly[3] for the execution of the following tasks: Paraphrase and reword, Grammar and spelling check. After using these tools, the authors reviewed and edited the content as needed. The authors assume complete responsibility for the content of the publication.

References

[1] European Union, Discover Europe's digital cultural heritage, 2026. www.europeana.eu/en.

[2] B. De Carolis, C. Gena, T. Kuflik, J. Lanir, Special issue on advanced interfaces for cultural heritage, International Journal of Human-Computer Studies 114 (2018) 1–2. doi:10.1016/j.ijhcs.2018.02.007.

[3] L. Ardissono, D. Petrelli, T. Kuflik, Personalization in cultural heritage: The road travelled and the one ahead, User Modeling and User-Adapted Interaction. Special Issue on Coming of Age: Celebrating a Quarter Century of User Modeling and Personalization 22 (2012) 73–99. doi:https://doi.org/10.1007/s11257-011-9104-x.

[4] E. Daga, R. Damiano, L. Díaz-Kommonen, T. Kuflik, M. Striani, Introduction to the special issue on extended intelligence for cultural engagement, J. Comput. Cult. Herit. 17 (2024). URL: https://doi.org/10.1145/3706102. doi:10.1145/3706102.

[5] J. a. Vilar, A. Rodrigues, N. Correia, An extended reality platform for cultural gaming: Enabling interactive narratives in spatial contexts, J. Comput. Cult. Herit. 17 (2024). URL: https://doi.org/10.1145/3702484. doi:10.1145/3702484.

[6] G. Volkmar, N. Wenig, R. Malaka, Memorial quest - a location-based serious game for cultural heritage preservation, in: Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts, CHI PLAY '18 Extended Abstracts, Association for Computing Machinery, New York, NY, USA, 2018, p. 661–668. URL: https://doi.org/10.1145/3270316.3271517. doi:10.1145/3270316.3271517.

[7] S. Varitimiadis, K. Kotis, D. Pittou, G. Konstantakis, Graph-based conversational AI: Towards a distributed and collaborative multi-chatbot approach for museums, Applied Sciences 11 (2021). URL: https://www.mdpi.com/2076-3417/11/19/9160. doi:10.3390/app11199160.

[8] M. Ribeiro, J. Santos, J. a. Lobo, S. Araújo, L. Magalhães, T. Adão, VR, AR, gamification and AI towards the next generation of systems supporting cultural heritage: addressing challenges of a museum context, in: Proceedings of the 29th International ACM Conference on 3D Web Technology, Web3D '24, Association for Computing Machinery, New York, NY, USA, 2024. URL: https://doi.org/10.1145/3665318.3677172. doi:10.1145/3665318.3677172.

[9] A. Chao, W. Pan, From data to narrative: A generative AI pipeline for cultural heritage audio guides, in: 2025 IEEE International Conference on Consumer Electronics - Taiwan (ICCE-Taiwan), 2025, pp. 771–772. doi:10.1109/ICCE-Taiwan66881.2025.11207912.

[10] A. Geninatti Cossatin, N. Mauro, F. Ferrero, L. Ardissono, Tell me more: integrating LLMs in a cultural heritage website for advanced information exploration support, Information Technology & Tourism (2025). doi:10.1007/s40558-025-00312-8.

[11] G. Trichopoulos, M. Konstantakis, G. Alexandridis, G. Caridakis, Large Language Models as recommendation systems in museums, Electronics 12 (2023). URL: https://www.mdpi.com/2079-9292/12/18/3829. doi:10.3390/electronics12183829.

[12] G. Trichopoulos, M. Konstantakis, G. Caridakis, A. Katifori, M. Koukouli, Crafting a museum guide using ChatGPT4, Big Data and Cognitive Computing 7 (2023). URL: https://www.mdpi.com/2504-2289/7/3/148. doi:10.3390/bdcc7030148.

[1] https://chatgpt.com/

[2] https://claude.ai/

[3] https://app.grammarly.com/

[13] S. Schauer, K. Simbeck, Conceptual implementation of a digital cultural heritage application using generative AI, in: Proceedings of the 21st International Conference on Culture and Computer Science: From Humanism to Digital Humanities, KUI '24, Association for Computing Machinery, New York, NY, USA, 2025. URL: https://doi.org/10.1145/3719236.3719242. doi:10.1145/3719236.3719242.

[14] A. Ferracani, M. Bertini, P. Pala, G. Nannotti, F. Principi, G. Becchi, Personalized generative storytelling with AI-visual illustrations for the promotion of knowledge in cultural heritage tourism, in: Proceedings of the 6th Workshop on the AnalySis, Understanding and ProMotion of HeritAge Contents, SUMAC '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 28–32. URL: https://doi.org/10.1145/3689094.3689465. doi:10.1145/3689094.3689465.

[15] G. Alexandridis, A. Chrysanthi, G. Tsekouras, G. Caridakis, Personalized and content adaptive cultural heritage path recommendation: an application to the gournia and Çatalhöyük archaeological sites, User Modeling and User-Adapted Interaction 29 (2019) 201–238. doi:10.1007/s11257-019-09227-6.

[16] M. L. Nappi, M. Buono, C. Chivăran, R. M. Giusto, Models and tools for the digital organisation of knowledge: accessible and adaptive narratives for cultural heritage, Heritage Science 12 (2024). doi:10.1186/s40494-024-01219-z.

[17] E. Not, D. Petrelli, Blending customisation, context-awareness and adaptivity for personalised tangible interaction in cultural heritage, International Journal of Human-Computer Studies 114 (2018) 3–19. URL: https://www.sciencedirect.com/science/article/pii/S1071581918300016. doi:https://doi.org/10.1016/j.ijhcs.2018.01.001, advanced User Interfaces for Cultural Heritage.

[18] A. Antoniou, A. Theodoropoulos, A. Chaleplioglou, E. Roinioti, P. Dafiotis, G. Lepouras, P. Chalari, G. Gujt, I. Lamprou, C. Sousa, AI-enabled cultural experiences: A comparison of narrative creation across different AI models, Electronics 14 (2025). URL: https://www.mdpi.com/2079-9292/14/20/4043. doi:10.3390/electronics14204043.

[19] A. Ferrato, C. Limongelli, F. Gasparetti, G. Sansonetti, A. Micarelli, Exploring the potential of multimodal large language models for question answering on artworks, in: Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '25, Association for Computing Machinery, New York, NY, USA, 2025, p. 432–436. URL: https://doi.org/10.1145/3708319.3733648. doi:10.1145/3708319.3733648.

[20] N. Xu, Y. Liu, Y. Chen, Z. Jiang, Y. Ren, Z. Zhang, B. Jia, L. Yu, Artifactshow: Incorporating generative AI into narrative visualization for interactive cultural experience, in: Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, CHI EA '25, Association for Computing Machinery, New York, NY, USA, 2025. URL: https://doi.org/10.1145/3706599.3720160. doi:10.1145/3706599.3720160.

[21] D. E. Tamah Al-Shammari, My heritage companion: An AI-driven mobile experience for visual storytelling, in: Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '25, Association for Computing Machinery, New York, NY, USA, 2025, p. 442–448. URL: https://doi.org/10.1145/3708319.3733651. doi:10.1145/3708319.3733651.

[22] O. Stock, M. Zancanaro, P. Busetta, C. Callaway, A. Krueger, M. Kruppa, T. Kuflik, E. Not, C. Rocchi, Adaptive, intelligent presentation of information for the museum visitor in PEACH, User Modeling and User-Adapted Interaction 17 (2007) 257–304. doi:https://doi.org/10.1007/s11257-007-9029-6.

[23] A. Zimmermann, A. Lorenz, LISTEN: a user-adaptive audio-augmented museum guide, User Modeling and User-Adapted Interaction 18 (2008) 389–416. doi:10.1007/s11257-008-9049-x.

[24] D. H. R. Spennemann, ChatGPT and the generation of digitally born "knowledge": How does a generative AI language model interpret cultural heritage values?, Knowledge 3 (2023) 480–512. URL: https://www.mdpi.com/2673-9585/3/3/32. doi:10.3390/knowledge3030032.

[25] T. D. True), N. Line, M. Mody, L. Hanks, J. Abbott, F. Acikgoz, A. Assaf, S. Bakir, A. Berbekova, A. Bilgihan, A. Dalton, E. Erkmen, M. Geronasso, D. Gomez, S. Graves, A. Iskender, S. Ivanov,

M. Kizildag, M. Lee, W. Lee, J. Luckett, S. McGinley, F. Okumus, I. Onder, O. Ozdemir, H. Park, A. Sharma, C. Suess, M. Uysal, T. Zhang, Generative artificial intelligence in the hospitality and tourism industry: Developing a framework for future research, Journal of Hospitality & Tourism Research 49 (2025) 235–253. URL: https://doi.org/10.1177/10963480231188663. doi:10.1177/10963480231188663.

[26] J. H. Kim, J. Kim, J. Park, C. Kim, J. Jhang, B. King, When ChatGPT gives incorrect answers: The impact of inaccurate information by Generative AI on tourism decision-making, Journal of Travel Research 0 (2025). URL: https://doi.org/10.1177/00472875231212996. doi:10.1177/00472875231212996.

[27] L. Mich, R. Garigliano, ChatGPT for e-tourism: a technological perspective, Information Technology & Tourism 25 (2023) 1–12. doi:10.1007/s40558-023-00248-x.

[28] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive NLP tasks, in: Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20, Curran Associates Inc., Red Hook, NY, USA, 2020. doi:10.5555/3495724.3496517.

[29] T. Nguyen, P.-M. Hui, F. M. Harper, L. Terveen, J. Konstan, Exploring the filter bubble: The effect of using recommender systems on content diversity, in: Proceedings of the 23rd International Conference on World Wide Web, WWW '14, Association for Computing Machinery, New York, NY, USA, 2014, p. 677–686. URL: https://doi.org/10.1145/2566486.2568012. doi:10.1145/2566486.2568012.

[30] S. Dai, C. Xu, S. Xu, L. Pang, Z. Dong, J. Xu, Bias and Unfairness in Information Retrieval Systems: New Challenges in the LLM Era, in: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24, Association for Computing Machinery, New York, NY, USA, 2024, pp. 6437–6447. URL: https://dl.acm.org/doi/10.1145/3637528.3671458. doi:10.1145/3637528.3671458.

[31] P. Kelly, J. Schild, A. Jafari, FolkRAG: a retrieval-augmented generation system for cultural heritage materials, Neural Computing and Applications 37 (2025) 20281–20297. URL: https://doi.org/10.1007/s00521-025-11455-4. doi:10.1007/s00521-025-11455-4.

[32] A. F. Chao, Blending Visitor Familiarity and Site Significance: RAG Enhanced Heritage Interpretation, in: 2024 IEEE 13th Global Conference on Consumer Electronics (GCCE), 2024, pp. 286–287. URL: https://ieeexplore.ieee.org/abstract/document/10760483. doi:10.1109/GCCE62371.2024.10760483, iSSN: 2693-0854.

[33] S. S. Y. Kim, Q. V. Liao, M. Vorvoreanu, S. Ballard, J. W. Vaughan, "I'm not sure, but...": Examining the impact of Large Language Models' uncertainty expression on user reliance and trust, in: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 822–835. URL: https://doi.org/10.1145/3630106.3658941. doi:10.1145/3630106.3658941.

[34] R. Dror, D. Hutchinson, M. Jones, V. Van Hyning, T. Kuflik, The curator's helper, in: Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 496–504. URL: https://doi.org/10.1145/3631700.3664905. doi:10.1145/3631700.3664905.

[35] L. Malmqvist, Sycophancy in Large Language Models: Causes and Mitigations, in: K. Arai (Ed.), Intelligent Computing, Springer Nature Switzerland, Cham, 2025, pp. 61–74. doi:10.1007/978-3-031-92611-2_5.

[36] Y. Daryani, Z. Sourati, M. Dehghani, The Homogenizing Engine: AI's Role in Standardizing Culture and the Path to Policy, Policy Insights from the Behavioral and Brain Sciences 13 (2026) 14–27. URL: https://doi.org/10.1177/23727322251406591. doi:10.1177/23727322251406591, publisher: SAGE Publications.