



From Interaction to Reflection: How AI Persona and Post-Experience Summarization Jointly Shape Aesthetic Anxiety and Interpretive Confidence in Art Exhibitions

Guanrong Lian

Zhejiang University, Hangzhou, China

3230104265@zju.edu.cn

Abstract. Many art-exhibition visitors experience aesthetic anxiety—a fear of misinterpreting artworks or being judged for unconventional responses. AI conversational systems offer new support, but two mechanisms—interaction style and post-experience summarization—have been studied in isolation. We report an exploratory 2×2 between-subjects experiment (N = 44) crossing AI interaction style (tool- vs. companion-type) with summarization in an art exhibition, triangulated with interviews. Interaction style showed no direct effects, but summarization showed a medium-to-large effect on retrospective interpretive confidence ($d = 0.63$, $p = .057$) benefiting primarily the companion condition ($p = .071$). Interviews attributed low expression pressure to the AI’s non-human nature and pointed to conversation depth—not the hypothesized cognitive openness—as the likely link from interaction to reflection. We discuss implications for museum AI balancing authority with non-judgmental support.

Keywords: Aesthetic anxiety, AI companion, Museum experience, Conversational AI, Post-experience reflection, Interpretive confidence

1 Introduction

Art exhibitions promise cultural enrichment, yet for many visitors an artwork produces not curiosity but *aesthetic anxiety*: a fear of misinterpreting what one sees and of being judged for an unconventional or “wrong” reading [1,2]. Facing the implicit authority of institutional spaces, visitors often choose silence over engagement. Conventional audio guides and panels only deepen this, operating one-directionally—authoritative knowledge flowing from institution to visitor—and thereby positioning understanding as something delivered correctly rather than constructed personally.

LLM-powered systems can instead converse adaptively. Two orientations have been studied separately: tool-type agents prioritizing informational accuracy and neutral authority, and companion-type agents that respond non-judgmentally, building on what the visitor expresses [9,11]. Art interpretation—where knowledge carries institutional authority but understanding is personal—sits between them and is under-examined. AI can also generate post-experience summaries [3,4], but this is treated as independent of

the preceding interaction. This exposes two gaps: which interaction design serves a context needing both authority and non-judgment, and whether conditions set during interaction moderate how well a later summary supports reflection. We propose the two are interdependent—a companion-type interaction that lowers defensiveness may create the openness for a summary to be integrated rather than received as further correction—and test this in a 2×2 experiment crossing interaction style with summarization.

Contributions. (1) A system: an LLM-based conversational system for art visitors with two configurable personas (tool- and companion-type) over a shared model, an artwork-specific knowledge base, and an automated summarization module. (2) An exploratory 2×2 experiment (N = 44) with interviews isolating the effects of interaction style, summarization, and their interaction. (3) An insight: psychological safety derived from the AI’s non-human nature rather than warmth, and conversation depth—not the hypothesized cognitive openness—as the likely link from interaction to reflection.

2 Related Work

Aesthetic anxiety. Art encounters are not uniformly pleasurable: schema violation triggers anxiety, and viewers either retreat or move through it via metacognitive reflection [1]. Coping potential—belief in one’s capacity to understand—decides whether complexity yields anxiety or curiosity [2], and can be raised by scaffolding, so an AI that builds on rather than corrects interpretation may help [5]. Yet classic accounts [6,7,8] warn that simply supplying “correct” readings can undermine authentic experience.

AI persona and reflection. Communication style shapes how users experience AI: companion-style personas reduce defensiveness and increase willingness to share uncertain responses [9,10], functioning as a low-pressure “sandbox” for judgment-free expression [11]—precisely what aesthetic anxiety forecloses. But such work assumes no correct answers, and anthropomorphic AI risks distortion when users over-identify with its framing [12]. Separately, AI summarization can reshape autobiographical narrative and improve articulation when it scaffolds rather than dictates [3,4]; when an AI’s output diverges from its assigned relational role, friction produces disengagement [13]. A summary after tool-type interaction thus arrives where the AI is already an authority; after companion-type, where the visitor’s interpretations have been validated. No study has examined this temporal relationship—our central gap.

3 Methodology

3.1 Study Design

We use an exploratory mixed-methods design: a 2×2 **between-subjects experiment** with post-session semi-structured interviews. Factors are **AI interaction style** (tool vs. companion) and **post-experience summarization** (present vs. absent), yielding groups A (tool, no summary), B (tool, summary), C (companion, summary), D (companion, no

summary). The study isolates the AI mechanisms from a prior robot-dog companion project to control for physical-form confounds.

3.2 Participants

Forty-four participants were recruited (23 female, 21 male; ages 18–35), excluding those with professional backgrounds in art, design, or art history [2]. Art familiarity, general self-efficacy, and baseline aesthetic anxiety were covariates. A balanced staged-adaptive allocation kept group sizes comparable: A ($n = 11$), B ($n = 10$), C ($n = 13$), D ($n = 10$).

3.3 Stimulus Materials

Artwork Selection. Four artworks (Table 1; Figure 1) were chosen for no dominant authoritative reading, emotional salience with open meaning, and diversity of interpretive challenge. An identical knowledge base per artwork (artist biography, context, statements, formal description) was provided to both AI conditions, so differences reflect interaction style, not content access.

Table 1. Selected artworks and interpretive challenge types

Artist	Work	Year	Challenge Type
Francis Bacon	Study after Velázquez's Portrait of Pope Innocent X	1953	Affective impact
Mark Rothko	No. 61 (Rust and Blue)	1953	Sensory immersion
Gerhard Richter	Abstraktes Bild	c.1997	Formal exploration
Cy Twombly	Untitled (Bolsena)	1969	Meaning suspension




(a) Bacon, 1953



(b) Rothko, 1953



(c) Richter, c. 1997



(d) Twombly, 1969

Fig. 1. Artworks presented to participants. © The Estate of Francis Bacon; © 1998 Kate Rothko Prizel & Christopher Rothko; © Gerhard Richter; © Cy Twombly Foundation.

AI System Design. Tool-type AI provided historical context, artist intentions, and mainstream interpretations in a clear, informative tone without soliciting feelings—the one-directional audio-guide model. Companion-type AI first validated the visitor’s initial response, then introduced context that deepened rather than replaced it, followed by open-ended questions, and never implied a single correct interpretation. Both used the same LLM, differing only by system prompt. Both used a one-time AI-initiated opening per artwork, holding initiation frequency constant: the tool agent opened directly (“Would you like to know more about this work?”); the companion agent, if no input came within 20 s, invited reflection (“This work seems to have caught your attention—would you like to share what you’re noticing?”). All later turns were visitor-driven.

Post-Experience Summary. Groups B and C received a personalized AI-generated summary after all four interactions, built from the full transcript. The prompt gave greater narrative emphasis to more-engaged artworks and applied a low-engagement safeguard (objective description only) when turns were few. Structured as a personal narrative—recurring themes, curatorial context anchored to the participant’s own observations, and a closing reflection—it was designed to mirror their experience rather than deliver a verdict.

3.4 Procedure

T1 (~3 min): baseline questionnaire (art familiarity, General Self-Efficacy Scale [14], aesthetic anxiety baseline). T2 (~10–15 min): sequential viewing of four artworks with AI chat; T2 questionnaire after all artworks but before any summary (expression pressure, engagement intention, in-situ confidence, cognitive openness, manipulation check). T3 (~5 min): summary groups read their summary; all completed the T3 questionnaire (retrospective confidence, structured sense-making, satisfaction), then a ~10 min interview. Total: ~18–23 min.

3.5 Measures and Analysis

All items used 7-point scales. T1: art familiarity (3; covariate), General Self-Efficacy [14] (10; covariate), aesthetic anxiety baseline (4; [2]). T2: expression pressure (4), engagement intention (3; [9]), in-situ confidence (3), cognitive openness (4; [10]; mediator), manipulation check (2). T3: retrospective confidence (3; [4]), structured sense-making (4), satisfaction (3); full instrument in Appendix A. Both personas share one model (DeepSeek-V4-flash), differing only by system prompt, with a knowledge base injected in full; B and C received a transcript-based summary at T3. Main effects (H1, H2) used one-way ANCOVAs with the three covariates, H3a a two-way ANCOVA (Type III), and mediation (H3b) a 5,000-resample bootstrap [15]. We report partial η^2 and Cohen’s d ($d \geq 0.5$ meaningful); interviews via reflexive thematic analysis [16].

4 RESULTS

We report quantitative findings then qualitative themes. Given the exploratory design and small sample ($N = 44$; cells 10–13), the study is powered only for large effects, and the moderated mediation has very limited power. We therefore treat effect sizes as primary and p -values as secondary descriptive indicators. All ANCOVAs control for art familiarity, general self-efficacy, and baseline aesthetic anxiety. No corrections for multiple comparisons were applied, reinforcing the emphasis on effect sizes. Interviews were in Chinese; quotes translated by the author.

4.1 Sample and Manipulation Check

No data were missing. One-way ANOVAs confirmed no significant baseline differences across groups on any covariate (all $p > .09$). The manipulation check ran in the expected directions but did not reach significance: companion-type rated higher on friendliness ($M = 5.35$ vs. 4.76 , $t(42) = 1.38$, $p = .175$, $d = 0.42$) and tool-type higher on sense of a “correct answer” ($M = 3.19$ vs. 2.78 , $p = .427$, $d = 0.24$). The stylistic contrast was perceived in the right direction but was not strong enough to fully separate conditions—a limitation we revisit.

4.2 H1: Main Effect of AI Interaction Style

Contrary to H1, interaction style had no significant main effect on any T2 outcome: expression pressure ($F(1,39) = 0.03$, $p = .860$), engagement intention ($p = .656$), in-situ confidence ($p = .754$). H1 was not supported. As the qualitative analysis shows, this null is informative: expression pressure was uniformly low because non-judgment was attributed to the AI’s non-human nature rather than its style.

4.3 H2: Main Effect of Post-Experience Summarization

Summarization had no effect on structured sense-making ($F(1,39) = 0.03$, $p = .860$; H2a not supported). For retrospective interpretive confidence, the summary condition ($M = 5.43$) exceeded no-summary ($M = 4.87$), $F(1,39) = 3.85$, $p = .057$, $\eta^2p = .090$, $d = 0.63$ —a medium-to-large trend just short of .05; H2b partially supported.

4.4 H3a: Interaction Between Style and Summarization

A two-way ANCOVA on retrospective confidence revealed a medium-sized interaction matching the hypothesis: $F(1,37) = 3.45$, $p = .071$, $\eta^2p = .085$. As Figure 2 shows, summarization helped primarily the companion condition ($4.74 \rightarrow 5.74$) while making virtually no difference for tool-type ($5.00 \rightarrow 5.03$). H3a received trend-level support.




Fig. 2. Interaction effect of AI style and summarization on retrospective interpretive confidence.

4.5 H3b: Moderated Mediation via Cognitive Openness

The hypothesized moderated mediation (mediator = cognitive openness) was not supported. The path from style to openness was non-significant ($a = -0.30, p = .298$); though openness strongly predicted retrospective confidence ($b = 0.80, p < .001$), conditional indirect effects' CIs included zero and the index of moderated mediation was non-significant ($0.03, 95\% \text{ CI } [-0.12, 0.33]$). H3b was not supported—the chain broke at path a, consistent with the weak manipulation.

4.6 Exploratory Analysis: Conversation Depth

Interviews suggested an alternative mechanism: companion-type interaction yields more and deeper conversation, hence a more fitting summary. Companion participants showed higher conversation volume, with total turns reaching a medium effect ($M = 19.6$ vs. $13.5, p = .078, d = 0.55$). Volume correlated with outcomes: user word count with retrospective confidence ($r = .48, p = .001$), and total turns with retrospective confidence ($r = .43$), openness ($r = .33$), and sense-making ($r = .37$). Among summary groups, companion (C) exceeded tool (B) on volume ($d = 0.64$) and retrospective confidence ($d = 0.70$). This gives convergent—if not conclusive—evidence for a conversation-depth pathway.

5 Discussion

The quantitative picture is coherent if tentative: the two core hypotheses (H2b, H3a) showed medium-to-large effects ($d \approx 0.61\text{--}0.63$) in the predicted direction, while direct style effects (H1) and the hypothesized mediation (H3b) were unsupported. The interviews illuminate why.

5.1 Why Style Had No Direct Effect.

Participants in both conditions reported low expression pressure, attributing it not to the AI's warmth but to its non-humanness. Psychological safety appears conferred by the absence of a human evaluator—a floor both styles share—so the companion style adds little incremental safety; for an anxiety rooted in social judgment, any AI interlocutor may lower the barrier simply by not being human. This is compounded by a ceiling: participants readily recognized the AI (“after all, it’s an AI,” C4), capping how friend-like the companion style could feel and weakening perception-based manipulation checks—even as a clear behavioral preference for being guided remained (C8). Users need not mistake the AI for a friend to benefit; the value is behavioral, not an illusion of friendship.

5.2 The Conversation-Depth Mechanism.

The hypothesized mechanism was unsupported; interviews instead point to a chain where companion-style interaction elicits more conversation, deepening engagement and giving the AI richer material for a fitting summary (“Its description felt like my own insight, now articulated back to me,” C6). This explains the H3a interaction: companion conversations were richer, while tool conversations stalled—one participant noted the AI “only answered what I asked, then wrapped up in three or four turns” (B3)—leaving thinner material. This is an exploratory, post-hoc account; its central claim, that richer dialogue yields a more fitting summary, was not directly measured and is offered as a hypothesis for pre-registered work.

5.3 Limitations.

Several limitations bound our conclusions. The manipulation was weak, so findings hold only under a limited stylistic contrast. The sample was small ($N = 44$), leaving key effects at trend level despite meaningful effect sizes. The central mechanism was not directly measured. There is also a domain limit: aesthetic experience is often pre-verbal and associative, whereas LLMs operate through propositional language, so even a well-designed companion AI may fail to meet the visitor in the intuitive register where meaning is first felt (B5). One baseline item was inadvertently omitted, and participants came from the author’s network, limiting generalizability.

6 Conclusion

We built an LLM-based conversational system for art exhibitions—two configurable personas and an automated summarization module over a shared model—to study how interaction style and summarization jointly shape aesthetic experience. In an exploratory 2×2 experiment ($N = 44$) with interviews, interaction style produced no direct effects, but summarization gave a medium-to-large benefit for retrospective interpretive

confidence, emerging primarily under companion-type interaction. Qualitatively, psychological safety derived from the AI's non-human nature rather than its warmth, and interaction appears linked to reflection through conversation depth rather than the hypothesized cognitive openness. The design implication is that interaction and reflection should be designed together—a summary is only as good as the conversation feeding it, and the value of companion-style AI lies in behavior (lowering the threshold to speak, sustaining dialogue) rather than simulated friendship. The two temporal phases appear interdependent; designing them in concert is the more promising path for museum AI. We offer the conversation-depth mechanism as a hypothesis for future pre-registered work.

References

1. Matthew Pelowski and Fuminori Akiba. 2011. A model of art perception, evaluation and emotion in transformative aesthetic experience. **New Ideas in Psychology** 29, 2, 80–110. <https://doi.org/10.1016/j.newideapsych.2010.04.001>.
2. Paul J. Silvia. 2005. Emotional responses to art: From collation and arousal to cognition and emotion. **Review of General Psychology** 9, 4, 342–357. <https://doi.org/10.1037/1089-2680.9.4.342>.
3. Lucy Osler. 2025. Knowing oneself through and with AI: From self-tracking to chatbots. arXiv:2512.03682. <https://arxiv.org/abs/2512.03682>.
4. Shou-Tzu Han. 2025. Narrative-Centered Emotional Reflection: Scaffolding Autonomous Emotional Literacy with AI. arXiv:2504.20342. <https://arxiv.org/abs/2504.20342>.
5. Helmut Leder, Benno Belke, Andries Oeberst, and Dorothee Augustin. 2004. A model of aesthetic appreciation and aesthetic judgments. **British Journal of Psychology** 95, 4, 489–508. <https://doi.org/10.1348/0007126042369811>.
6. Pierre Bourdieu. 1984. **Distinction: A Social Critique of the Judgement of Taste**. Harvard University Press. <https://www.hup.harvard.edu/catalog.php?isbn=9780674212770>.
7. John Dewey. 1934. **Art as Experience**. Minton, Balch & Company. <https://archive.org/details/in.ernet.dli.2015.188318>.
8. James O. Young. 2001. **Art and Knowledge**. Routledge. <https://www.routledge.com/Art-and-Knowledge/Young/p/book/9780415256472>.
9. Juhye Ha, Hyeon Jeon, Daeun Han, Jinwook Seo, and Changhoon Oh. 2024. CloChat: Understanding How People Customize, Interact, and Experience Personas in LLMs. arXiv:2402.15265. <https://arxiv.org/abs/2402.15265>.
10. Andreas Janson. 2023. How to leverage anthropomorphism for chatbot service interfaces: The interplay of communication style and personification. **Computers in Human Behavior** 149, 107954. <https://doi.org/10.1016/j.chb.2023.107954>.
11. Renkai Ma et al. 2026. Negotiating Digital Identities with AI Companions. In **Proc. CHI '26**. ACM. arXiv:2601.12181. <https://arxiv.org/abs/2601.12181>.
12. Myra Cheng, Alicia DeVrio, Lisa Egede, Su Lin Blodgett, and Alexandra Olteanu. 2024. Understanding the Impact of GenAI Requires Understanding the Impact of Anthropomorphic AI. arXiv:2410.08526. <https://arxiv.org/abs/2410.08526>.
13. Emrecaan Gulay, Eleonora Picco, Enrico Glerean, and Corinna Coupette. 2026. Relational Dissonance in Human-AI Interactions. arXiv:2509.15836. <https://arxiv.org/abs/2509.15836>.

14. Ralf Schwarzer and Matthias Jerusalem. 1995. Generalized Self-Efficacy Scale. In **Measures in Health Psychology**, NFER-NELSON, 35–37. <http://userpage.fu-berlin.de/~health/engscal.htm>.
15. Andrew F. Hayes. 2018. **Introduction to Mediation, Moderation, and Conditional Process Analysis** (2nd ed.). Guilford Press. <https://www.guilford.com/books/Introduction-to-Mediation-Moderation-and-Conditional-Process-Analysis/Andrew-Hayes/9781462534654>.
16. Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. **Qualitative Research in Psychology** 3, 2, 77–101. <https://doi.org/10.1191/1478088706qp063oa>.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

