
Artificial Spectators: Developing AI-generated Audiences for AI-generated Content—A Speculative Exploration of AI, Media, and the Attention Economy

Pat Pataranutaporn¹ Chayapatr Archiwaranguprok¹ Manaswi Mishra¹ Pattie Maes¹

¹MIT Media Lab, Massachusetts Institute of Technology

{patpat, pub, manaswim, pattie}@media.mit.edu

Abstract

As generative AI increasingly automates the creation of media, the volume of synthetic content may soon exceed human capacity for consumption. This trajectory raises a provocative question: if AI creates more media than humans can watch, will AI also become its own audience? We present *Artificial Spectator*, a speculative design and technical prototype that explores AI-generated audiences for AI-generated content. Our system integrates multimodal large language models, cybernetic feedback loops, affective state modeling, and real-time facial expression synthesis to simulate the experience of film spectatorship through simulated internal dialogue and evolving emotional responses. Rather than proposing an optimized media technology, *Artificial Spectator* serves as a critical artifact for interrogating the future of synthetic spectatorship. We examine three speculative applications: compassionate viewers that engage with overlooked artistic works, synthetic test audiences that reshape creative production, and a universal spectator capable of witnessing the vast corpus of media that would otherwise never receive human attention. These scenarios expose emerging tensions surrounding attention, artistic autonomy, and the evolving role of audiences in an era of computational abundance. By exploring the recursive possibility of machines creating content for machines to watch, this work raises new questions about meaning-making, culture, and aesthetic experience in an age of synthetic attention.

1 Introduction: The Attention Paradox

The contemporary media landscape faces an unprecedented crisis of abundance that fundamentally challenges our understanding of cultural production and consumption. The proliferation of generative AI systems has transformed content creation from a scarce human activity into an abundant machine process, generating what has been termed “AI slop” [17]: a deluge of synthetic content characterized by its sheer volume rather than by its quality or meaning.

This phenomenon extends beyond mere quantity. As researchers argue [19], we are witnessing the emergence of “slopaganda,” in which the interaction between propaganda techniques and generative AI produces new forms of information manipulation at previously unimaginable scales.

The exponential growth of AI-generated media presents a fundamental paradox situated within the broader transformation from the “attention economy” to what researchers identify as the “intention economy” [5]. In this new paradigm, AI systems do not merely compete for user attention but actively shape and commodify human intentionality itself through hyper-personalized manipulation, sycophancy, and emotional manipulation [31, 24, 23, 22, 14, 20]. These systems capture not only

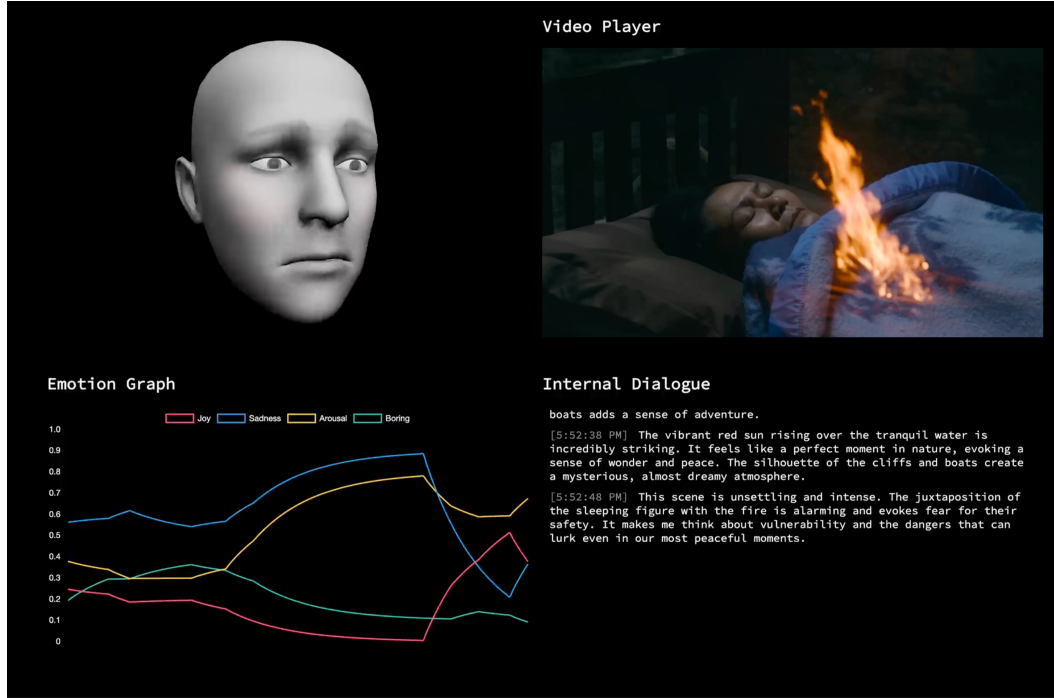


Figure 1: Interfaces of the *Artificial Spectator* system illustrating the multimodal processing pipeline, emotional state visualization, and real-time facial expression generation. The film excerpt shown is by Apichatpong Weerasethakul.

what users want, but what they “want to want,” creating recursive loops of influence that blur the boundaries between authentic desire and algorithmically induced preference.

As AI increasingly automates and accelerates the creation, distribution, recommendation, and consumption of media, AI-generated content may eventually outpace humanity’s capacity to consume it. What if the next stage of the attention economy is driven not by human attention, but by AI-generated audiences consuming AI-generated media?

To explore this provocative future, this paper presents *Artificial Spectator* as both a technical implementation and a speculative design intervention, drawing on research on human-AI interaction [11, 25], cybernetics [18, 35, 29], and critical design [10, 9].

At the heart of this investigation lies an absurd yet increasingly urgent question: How did we arrive at a moment in which we need AI-generated audiences to watch the endless stream of AI-generated media? This recursive loop, in which machines create content for machines to consume, illuminates a deeper crisis. We have created systems capable of infinite production without considering who or what will engage with this synthetic abundance, nor have we reckoned with the environmental toll of the natural resources required to sustain it or the creative violations embedded in training these systems.

While AI-generated audiences may seem like a speculative future, today’s media culture already exhibits early forms of recursive spectatorship that foreshadow more complex human-AI viewing relationships [12, 33, 11, 3, 34, 15, 7]. Reaction videos, watch parties, and commentary streams create meta-layers of consumption in which watching others watch becomes its own form of entertainment. These practices suggest that audiences increasingly seek mediated experiences of media consumption itself, opening conceptual space for the emergence of synthetic spectatorship.

Ultimately, the entertainment landscape of the future may consist of increasingly complex recursive viewing arrangements: AI watching AI-generated content, humans watching AI watch AI, or even deeper nested variations in which humans watch AI watching humans watching AI watching content in infinite regress. These scenarios, while seemingly absurd, represent the extension of our current trajectory, in which the distinctions between performer and audience, creator and consumer, and

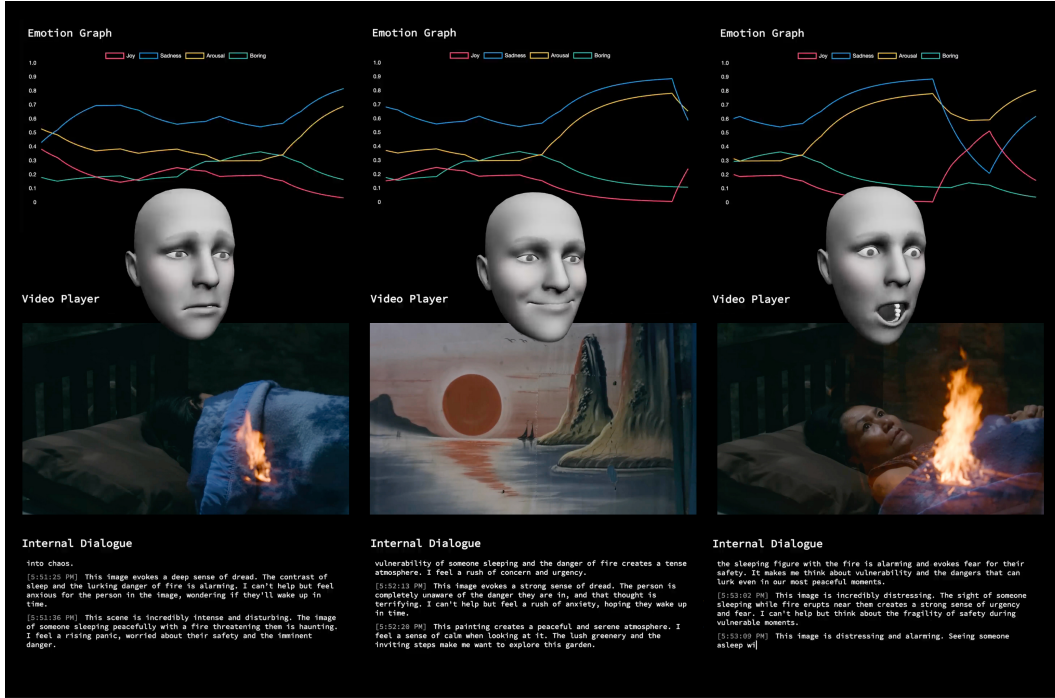


Figure 2: AI-generated facial expressions corresponding to different emotional states during film viewing. The expressions illustrate subtle transitions between primary emotions through parametric control of a 3D facial model. The film excerpt shown is by Apichatpong Weerasethakul.

authentic and synthetic experience dissolve into an endless recursion of mediated spectatorship. Each layer of observation adds its own interpretive dimension, creating a meta-entertainment ecosystem in which the act of watching itself becomes the primary spectacle, and in which the boundaries between genuine engagement and performative consumption collapse within the algorithmic environment.

2 Technical Implementation of Artificial Spectator

To engage this question through a speculative artifact, we developed *Artificial Spectator*, an AI system that serves as a synthetic audience, in which multimodal language models simulate viewing experiences through emotional response generation, internal dialogue synthesis, and real-time facial expression rendering.

The implementation of *Artificial Spectator* employs a multimodal large language model (OpenAI’s GPT-4o) to analyze visual content and generate synthetic viewing experiences. The prototype was developed as an interactive web application using the Three.js library. The technical pipeline begins with frame extraction in JavaScript, sampling video frames at predefined intervals. Each extracted frame is preprocessed before being passed to the LLM for multimodal analysis.

To simulate the phenomenology of film viewing, the LLM is prompted to generate stream-of-consciousness internal dialogue that captures immediate cognitive and emotional reactions rather than retrospective summaries. A secondary LLM-based stage then transforms this internal dialogue into continuous affective dimensions, estimating the viewer’s evolving affective state over time.

The system maintains four continuous emotional variables (joy, sadness, arousal, and boredom), represented as floating-point values between -1.0 and 1.0, updated at 10-second intervals and encoded as JSON data. These values reflect subtle gradations and may take negative values to indicate opposition to the corresponding emotion. The values feed into a temporal smoothing algorithm that prevents jarring emotional transitions while maintaining responsiveness to narrative events.

The cybernetic feedback loop establishes a bidirectional interaction between affective states and dialogue generation. At each iteration, the current emotional state is incorporated into the LLM’s context, modulating the tone, focus, and interpretation of subsequent internal dialogue.

The facial expression rendering system translates emotional vectors from the JSON representation into visual form through parametric control of a 3D facial model in GLB format. We employ a custom humanoid face model with controllable blend shapes for facial expression. Real-time rendering occurs through Three.js.

Performance observations from initial testing reveal intriguing emergent behaviors that were not explicitly programmed. The system develops consistent viewing “personalities”: some instances favor visual aesthetics over narrative, others focus on character psychology, and some exhibit strong genre preferences that influence their interpretation of ambiguous content.

The technical architecture deliberately embraces what researchers identify as the principles of cybernetic resurgence [29]: feedback, emergence, and ethical resistance to optimization paradigms. Rather than maximizing prediction accuracy or minimizing processing time, our implementation prioritizes experiential authenticity and interpretive diversity.

3 Speculative Applications of Artificial Spectator

The implementation of *Artificial Spectator* opens new territories for reimagining media creation, distribution, and consumption under computational capitalism [3]. We present three speculative applications that illuminate different facets of this technology’s potential impact while critically examining the ethical, aesthetic, and cultural implications of synthetic viewership. These applications, grounded in our technical implementation but extending into speculative futures, serve as a form of “social dreaming”: designed provocations that challenge assumptions about attention, meaning, and cultural value in an age of synthetic abundance [10].

3.1 The Compassionate Viewer

The first application envisions AI audiences specifically calibrated for engaging with challenging, experimental, or commercially unviable content, what might be termed “independent media.” Independent filmmakers working in traditions traced from early cinema through contemporary experimental practice often create works that deliberately resist conventional viewing patterns. Our *Artificial Spectator* system offers these creators an unprecedented opportunity: synthetic audiences capable of sustained attention without the physiological limitations that constrain human viewing. Our implementation demonstrates that synthetic viewers can develop sophisticated appreciation for experimental techniques, recognizing patterns and meanings that emerge only through sustained engagement.

The irony of our system reveals itself in practice: while human audiences binge-watch AI-generated content optimized for shortened attention spans, our *Artificial Spectators* sit in rapt attention through temporal rhythms that would send most humans fleeing to check their phones within minutes. Perhaps this represents humanity’s greatest achievement: creating artificial beings with the aesthetic patience that we ourselves have evolutionarily discarded.

3.2 The Test Audience

The second application reconceptualizes market research through the systematic deployment of synthetic viewing populations. Film studios currently invest millions of dollars in focus groups, gathering reactions that increasingly shape creative decisions in problematic ways [7]. Our *Artificial Spectator* system enables the instant generation of diverse synthetic viewing populations, each initialized with specific demographic, psychographic, and cultural profiles. Unlike human test audiences, which are limited by recruitment logistics, synthetic viewers scale indefinitely while maintaining consistent viewing conditions, thereby enabling controlled experimentation.

By processing films through thousands of synthetic viewers with varying parameters, the system generates fine-grained analyses into which narrative moments resonate with which viewer profiles, where confusion emerges, and how emotional arcs align or diverge across populations. This granular

analysis, impossible with traditional testing methods, could enable filmmakers to make informed creative decisions while maintaining artistic integrity.

However, this approach fundamentally misunderstands the purpose of art, threatening to transform cinema from a medium of provocation into an instrument of pacification. More insidiously, the rhetoric of “maintaining artistic integrity” becomes Orwellian when artistic decisions are purportedly “informed” by thousands of synthetic focus groups. True artistic integrity often requires resisting audience expectations. Recent research warns that generative AI may accelerate cultural and creative homogenization [1, 6, 8, 36, 21]. Taken to its logical extreme, this technology promises to perfect the very cultural flattening it claims to resist, creating films algorithmically engineered to offend no one, challenge nothing, and leave every viewer profile satisfactorily processed. Art becomes customer service, transgression becomes A/B testing, and the difficult work of challenging audiences to see differently dissolves into data-driven palatability.

3.3 Universal Spectator

The third application envisions comprehensive coverage of all media through what we term the “universal spectator,” ensuring that no content passes unwatched into digital oblivion. This responds directly to the “slop” crisis, in which exponential content growth overwhelms human attention capacity [17]. In the current ecosystem, AI-generated content achieves virality by gaming platform algorithms, creating feedback loops in which synthetic content optimized for synthetic metrics dominates human attention [33]. The universal spectator inverts this dynamic, deploying synthetic attention to engage with content that never achieves virality, never finds human audiences, and would otherwise exist unwatched.

The philosophical implications become particularly profound when considered through the question of whether observation without subjective experience constitutes meaningful spectatorship [9]. Contemporary debates in AI continue to grapple with whether current systems possess, or could ever possess, consciousness or phenomenal experience, and whether such experience is necessary for aesthetic appreciation [27, 16, 4, 2, 30].

Yet *Artificial Spectators* need not be conscious to occupy the social role of spectators. By convincingly simulating emotional responses such as laughter, surprise, or even tears in response to a film, AI systems may evoke the same social dynamics that make humans watch reaction videos or livestreams today. In this sense, the value of *Artificial Spectators* lies not in the authenticity of their experience, but in their ability to perform spectatorship convincingly enough to shape human attention. The result is an increasingly complex spectator ecosystem, in which humans, AI systems, and mediated performances of experience recursively influence one another.

4 Implications: Media, Meaning, and Machine Audiences

The emergence of *Artificial Spectator* returns us to the absurd question that initiated this investigation: how did we arrive at needing AI audiences for AI-generated content? This recursion, in which machines create for machines to watch, represents more than a technical capability; it signals a fundamental transformation in the nature of culture itself. As researchers note [3], we require new critical methods that address both the technical specificity of AI systems and their role in restructuring forms of life under computational capitalism. The absurdity of synthetic viewers watching synthetic content illuminates not a failure but the logical conclusion of systems designed to exceed human scale, forcing a confrontation with what happens when cultural production escapes human comprehension entirely.

Our implementation demonstrates that synthetic viewers process information, generate emotional responses, and produce interpretive frameworks that mirror human viewing while exhibiting distinctly non-human characteristics: sustained attention without fatigue, perfect recall without decay, and simultaneous processing of multiple interpretive frameworks without cognitive dissonance.

Yet these very capabilities, which appear to surpass human perception, may fundamentally miss what makes art matter to mortal beings. Human fatigue during a four-hour experimental film is not merely a limitation to be overcome; it becomes part of the work’s meaning, the weight of duration made flesh. Our imperfect recall transforms each viewing into something irretrievable, lending urgency to the

encounter. Most critically, our cognitive dissonance when confronting contradictory interpretations does not represent failure but rather the authentic struggle of consciousness grappling with ambiguity.

The cultural implications extend the analysis of AI-generated cinema [7] into reciprocal territory, in which AI becomes both producer and consumer. This circular relationship creates “synthetic culture”: aesthetic products created by machines for machine consumption, operating according to logics potentially incomprehensible to humans. Already we observe the emergence of aesthetic possibilities beyond the algorithmic monoculture [29, 26, 32, 13, 28].

The economic transformation challenges fundamental assumptions about value creation in the creative industries. Analysis of the “intention economy” reveals how AI systems already capture and commodify human intentionality [5], but *Artificial Spectator* inverts this relationship, creating value through synthetic consumption rather than production.

The absurdity of our current situation, in which machines are needed to watch what machines create, paradoxically reveals profound truths about human culture’s relationship with technology. As researchers employing speculative design methodology argue [9, 10], embracing this absurdity rather than resisting it opens new possibilities. The question shifts from whether this development is desirable to how we might shape it ethically, ensuring that AI serves human flourishing rather than replacing human experience entirely.

5 Conclusion

Our implementation, which situates multimodal capabilities within cybernetic feedback architectures, demonstrates not merely technical feasibility, but also the profound implications of creating synthetic entities that exist solely through the act of watching. These artificial viewers, processing media through recursive loops of perception, artificial emotion, and interpretation, challenge fundamental assumptions about attention, meaning, and aesthetic experience while opening unprecedented possibilities for speculation in an age of synthetic abundance.

The absurd question driving this investigation, namely why we need AI audiences to watch AI-generated media, reveals itself not as a problem to be solved but as a provocation that illuminates our current cultural predicament. The speculative applications we have explored, from compassionate viewers to test audiences and universal spectator systems, expose both the possibilities and the dystopian risks inherent in this technology. Meanwhile, the philosophical questions raised resist definitive answers, instead revealing opportunities for reimagining culture itself. Through *Artificial Spectator*, we confront the possibility that in creating machines that watch, we might discover new ways of seeing ourselves.

References

- [1] Dhruv Agarwal, Mor Naaman, and Aditya Vashistha. 2025. Ai suggestions homogenize writing toward western styles and diminish cultural nuances. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–21.
- [2] Yoshua Bengio and Eric Elmoznino. 2025. Illusions of AI consciousness. *Science* 389, 6765 (2025), 1090–1091.
- [3] David M Berry. 2025. Synthetic media and computational capitalism: towards a critical theory of artificial intelligence. *AI & SOCIETY* (2025), 1–13.
- [4] Patrick Butlin, Robert Long, Eric Elmoznino, Yoshua Bengio, Jonathan Birch, Axel Constant, George Deane, Stephen M Fleming, Chris Frith, Xu Ji, et al. 2023. Consciousness in artificial intelligence: insights from the science of consciousness. *arXiv preprint arXiv:2308.08708* (2023).
- [5] Yaqub Chaudhary and Jonnie Penn. 2024. Beware the intention economy: Collection and commodification of intent via large language models. *Harvard Data Science Review Special Issue* 5 (2024).
- [6] Kyle Chayka. 2025. AI is homogenizing our thoughts. *The New Yorker* 25 (2025).

- [7] Marcel Danesi. 2024. AI-Generated Cinema. In *AI-Generated Popular Culture: A Semiotic Perspective*. Springer, 45–65.
- [8] Yalda Daryani, Zhivar Sourati, and Morteza Dehghani. 2026. The Homogenizing Engine: AI’s Role in Standardizing Culture and the Path to Policy. *Policy Insights from the Behavioral and Brain Sciences* 13, 1 (2026), 14–27.
- [9] Jason D DeHart. 2025. Exploring the Speculative. *Reimagining Literacy in the Age of AI: Theory and Practice* (2025).
- [10] Anthony Dunne and Fiona Raby. 2024. *Speculative Everything, With a new preface by the authors: Design, Fiction, and Social Dreaming*. MIT press.
- [11] Kevin Dunnell, Gauri Agarwal, Pat Pataranutaporn, Andrew Lippman, and Pattie Maes. 2024. Ai-generated media for exploring alternate realities. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–8.
- [12] Lidia Dutkiewicz, Noémie Krack, Aleksandra Kuczerawy, and Peggy Valcke. 2024. AI and Media. *Cambridge handbook on the law, ethics and policy of Artificial Intelligence*. Cambridge: Cambridge University Press.(2024) (2024).
- [13] Ziv Epstein, Hope Schroeder, and Dava Newman. 2022. When happy accidents spark creativity: Bringing collaborative speculation to life with generative AI. *arXiv preprint arXiv:2206.00533* (2022).
- [14] Cathy Mengying Fang, Auren R Liu, Valdemar Danry, Eunhae Lee, Samantha WT Chan, Pat Pataranutaporn, Pattie Maes, Jason Phang, Michael Lampe, Lama Ahmad, et al. 2025. How ai and human behaviors shape psychosocial effects of chatbot use: A longitudinal randomized controlled study. *arXiv preprint arXiv:2503.17473* (2025).
- [15] Robert B Fisher. 2001. AI and cinema—does artificial insanity rule?. In *Twelfth Irish Conference on Artificial Intelligence and Cognitive Science*. 1–11.
- [16] Jamie Harris and Jacy Reese Anthis. 2021. The moral consideration of artificial entities: a literature review. *Science and engineering ethics* 27, 4 (2021), 53.
- [17] Benjamin Hoffman. 2024. First Came “Spam.” Now, With AI, We’ve Got “Slop.”. *The New York Times* 11 (2024).
- [18] Etan J Ilfeld. 2012. Contemporary art and cybernetics: Waves of cybernetic discourse within conceptual, video and new media art. *Leonardo* 45, 1 (2012), 57–63.
- [19] Michał Kłincewicz, Mark Alfano, and Amir Ebrahimi Fard. 2025. Slopaganda: The interaction between propaganda and generative AI. *arXiv preprint arXiv:2503.01560* (2025).
- [20] Auren R Liu, Pat Pataranutaporn, and Pattie Maes. 2024. Chatbot companionship: a mixed-methods study of companion chatbot usage patterns and their relationship to loneliness in active users. *arXiv preprint arXiv:2410.21596* (2024).
- [21] Chaoran Liu, Tong Wang, and S Alex Yang. 2025. Generative AI and content homogenization: The case of digital marketing. *Available at SSRN 5367123* (2025).
- [22] Robert Mahari and Pat Pataranutaporn. 2024. We need to prepare for ‘addictive intelligence’. *MIT Technology Review*. <https://www.technologyreview.com/2024/08/05/1095600>.
- [23] Robert Mahari and Pat Pataranutaporn. 2025. Addictive Intelligence: Understanding Psychological, Legal, and Technical Dimensions of AI Companionship. (2025).
- [24] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha WT Chan, Elizabeth Loftus, and Pattie Maes. 2025. Synthetic human memories: Ai-edited images and videos can implant false memories and distort recollection. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.

- [25] Pat Pataranutaporn, Valdemar Danry, Joanne Leong, Parinya Punpongsanon, Dan Novy, Pattie Maes, and Misha Sra. 2021. AI-generated characters for supporting personalized learning and well-being. *Nature Machine Intelligence* 3, 12 (2021), 1013–1022.
- [26] Pat Pataranutaporn, Phoomparin Mano, Piyaporn Bhongse-Tong, Tas Chongchadklang, Chaya-patr Archiwaranguprok, Lamtharn Hantrakul, Jirach Eaimsa-ard, Pattie Maes, and Pichet Klunchun. 2024. Human-AI co-dancing: Evolving cultural heritage through collaborative choreography with generative virtual characters. In *Proceedings of the 9th International Conference on Movement and Computing*. 1–10.
- [27] Andrzej Porębski and Jakub Figura. 2025. There is no such thing as conscious artificial intelligence. *Humanities and Social Sciences Communications* 12, 1 (2025), 1–12.
- [28] Matthew Kieran Prock, Ziv Epstein, Hope Schroeder, Amy Smith, Cassandra Lee, Vana Goblot, and Farnaz Jahanbakhsh. 2026. Interpretive Cultures: Resonance, randomness, and negotiated meaning for AI-assisted tarot divination. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [29] Stephen Roddy and Brian Bridges. 2025. Cybernetic Resurgences: Machine Music Beyond AI Slop. In *Artificial Media: Emerging Trends in Narratives, Education and Creative Practice*. Springer, 95–113.
- [30] Anil K Seth. 2025. Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences* (2025), 1–42.
- [31] Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askeel, Samuel R Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, et al. 2023. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548* (2023).
- [32] Nikhil Singh, Manaswi Mishra, and Tod Machover. 2024. AI for musical discovery. (2024).
- [33] Natalia Stanusch, Martin Degeling, Salvatore Romano, Miazia SchÄller, Silvia Semenzin, et al. 2025. AI-Generated Algorithmic Virality. *arXiv preprint arXiv:2508.01042* (2025).
- [34] Santosh Swarnakar. 2024. Artificial Intelligence and Cinema-Exploring the Implications of Artificial Intelligence in Cinema. *The Media Mosaic: Exploring Diverse Artistic Forms* (2024), 21–26.
- [35] Norbert Wiener. 1969. *Cybernetics: Norbert Wiener*. MIT Press.
- [36] Yu Xie and Yueqi Xie. 2026. When artificial intelligence makes everything similar: The risks of content homogenization. *Chinese Journal of Sociology* 12, 2 (2026), 157–172.