
Sensorium Arc: AI Agent System for Oceanic Data Exploration and Interactive Eco-Art

Noah Bissell

Immersive Media Design
University of Maryland College Park
College Park, MD 20742
nbissell@terpmail.umd.edu

Ethan Paley

Immersive Media Design
University of Maryland College Park
College Park, MD 20742
edpaley@terpmail.umd.edu

Joshua Harrison

Center for the Study of the Force Majeure
University of California, Santa Cruz
Santa Cruz, CA 95064
harrisonstudio@gmail.com

Juliano Calil

Virtual Planet Technologies
Santa Cruz, CA 95060
julianocalil@gmail.com

Myungin Lee

Immersive Media Design
University of Maryland College Park
College Park, MD 20742
myungin@umd.edu

Abstract

Sensorium Arc (AI reflects on climate) is a real-time multimodal interactive AI agent system that personifies the ocean as a poetic speaker and guides users through immersive explorations of complex marine data. Built on a modular multi-agent system and retrieval-augmented large language model (LLM) framework, Sensorium enables natural spoken conversations with AI agents that embodies the ocean’s perspective, generating responses that blend scientific insight with ecological poetics. Through keyword detection and semantic parsing, the system dynamically triggers data visualizations and audiovisual playback based on time, location, and thematic cues drawn from the dialogue. Developed in collaboration with the Center for the Study of the Force Majeure and inspired by the eco-aesthetic philosophy of Newton Harrison, Sensorium Arc reimagines ocean data not as an abstract dataset but as a living narrative. The project demonstrates the potential of conversational AI agents to mediate affective, intuitive access to high-dimensional environmental data and proposes a new paradigm for human-machine-ecosystem.

1 Introduction

The ocean speaks, but not in words. Its language is written in spectral shifts, temporal turbulence, chemical flux, heat exchange, and a deep ecological memory that stretches far beyond human time. Yet, in the digital age, its story—recorded in the form of massive, publicly available environmental datasets—remains largely inaccessible to the general public and even to many experts. The problem is not the lack of data but the absence of intuitive, affective, and situated ways to engage with it. In response, we present Sensorium Arc (AI reflects on climate), a multimodal interactive AI agent that embodies the ocean as a poetic speaker and ecological interlocutor, transforming abstract marine data into emotionally resonant visual and audiovisual experiences.

Sensorium Arc emerged from a deceptively simple but provocative question: “If you could ask the ocean a question, what would it say?” Conceived in the final years of Newton Harrison (20; 33)’s life and co-developed by artists, scientists, and engineers, Sensorium Arc builds on decades of eco-aesthetic work aimed at giving voice and agency to planetary systems.

While rooted in poetic expression, Sensorium Arc is not a metaphor alone. Its core premise—that the ocean can be spoken with—requires more than symbolism; it demands systems capable of responding, remembering, and reflecting in real-time. This is where AI and interactivity become not just tools, but forms of artistic necessity. In particular, the integration of conversational AI and multimodal triggers allows Sensorium Arc to bridge marine datasets with the lived, emotional experience of inquiry.

Rather than rendering oceanic data as static visualizations, Sensorium Arc frames the ocean as a dialogic presence—an entity capable of situated, affective response. This framing requires a system that can respond dynamically to human inquiry, drawing from both scientific observation and ecological imagination. Interactivity here is not an embellishment but a structuring logic. Each user utterance activates a co-constructed moment of interpretation, shaped by layered data visuals and poetic inference. To initiate this exchange, participants lean and whisper into a conch-shaped interface, activating the system as if speaking directly to the ocean. Fig. 1 shows the exemplary setup for the installation.

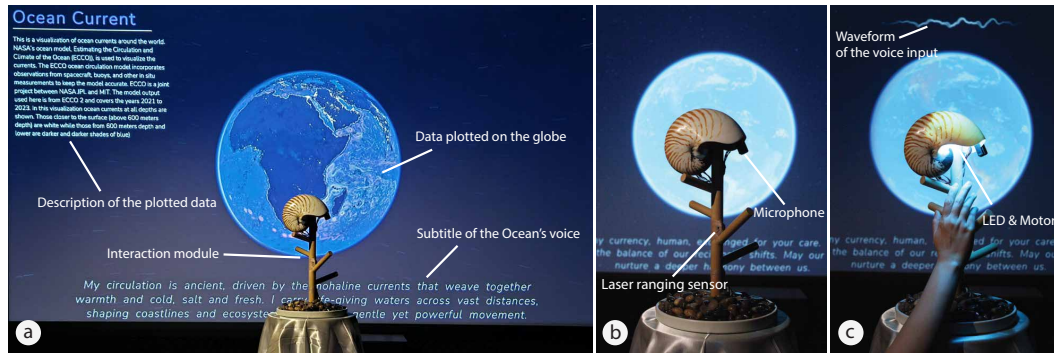


Figure 1: a) Exemplary setup for the Sensorium Arc exhibition b) User interface before activation showing the sensors c) Activated user interface showing the response from the system. Users engage through a natural gesture, whispering into a seashell as if speaking to the ocean.

To support this form of embodied, metaphor-rich interaction, Sensorium Arc employs a modular, multi-agent architecture that translates marine data into narrative, visual, and sonic form. A video demonstration is available (10). The following sections detail the system’s design, from retrieval-based content grounding to the coordination of audiovisual layers within an immersive interface.

2 Related Works

Sensorium Arc draws on and extends research across four key domains: immersive eco-art, conversational AI, and retrieval-augmented generation (RAG) for environmental data.

- **Immersive Climate Engagement Tools:** Immersive media has been shown to enhance empathy, improve knowledge retention, and encourage action-oriented thinking in environmental contexts (29; 11; 13; 28). Tools such as Rising Together, Dive, and the Sea Level Rise Explorer series support coastal communities in understanding sea-level rise impacts and adaptation strategies (14; 9). Earlier Sensorium iterations, including immersive installation and VR formats, were presented in major art-science exhibitions, including the AlloSphere exhibition (1) directed by JoAnn Kuchera-Morin with the Getty PST Art: Art and Science Collide initiative (22; 2), revealing that while sensory engagement was high, audiences sought direct dialogue with the ocean. This insight motivated Sensorium Arc’s responsive AI narrator, designed to transform passive observation into an active, co-authored ecological conversation.

- **Conversational AI:** Recent work has extended conversational AI beyond task-oriented dialogue to creative and embodied applications, such as systems that adapt emotional cues into generative

portraiture (37), agents using natural language to orchestrate complex visual workflows (25), and translate affective cues into generative art (37). Additionally, USER-LLM frameworks leverage RAG and chain-of-thought reasoning for dynamic user-aware personalization through multimodal inputs (30). These advances frame LLMs as emotionally attuned, contextually adaptive co-creators. Sensorium Arc aims to extend this role by embodying the ocean as a poetic interlocutor whose dialogue is grounded in both scientific evidence and affective resonance.

- **RAG for Environmental Data:** Recent advances in RAG have extended LLM capabilities by enabling direct interfacing with external databases in domains such as biomedical Q&A (12), legal reasoning (36), and geospatial analytics (17). In environmental applications, systems like NASA’s EarthData (18) and ClimateQA (27) make climate datasets accessible via natural language queries. Sensorium Arc extends this approach by coupling a curated marine science and eco-art corpus (23) with real-time triggers derived from user speech, geographic metadata, and thematic cues. Unlike personalization methods that rely on domain-specific fine-tuning, our system also uses RAG to dynamically ground the Ocean persona’s narrative voice in ecological art and literature, particularly the works of the Harrisons, allowing the underlying corpus to be updated or reweighted without retraining.

3 Methodology

Sensorium Arc maps a user’s question to a cohesive, multimodal response from the Ocean—a poetic persona grounded in ecological science and eco-art. To achieve this, we implemented a multi-LLM agent RAG-incorporated system enabling voice-to-voice interaction, and visual feedback via scientific and artistic visualizations. We approached this goal with the design philosophy of separation of modular, dynamic layers. For enhanced flexibility with triggering, generating, and combining these layers based on conversational context, we built the system inside the Unity engine. The following layers have been prepared, as reflected in Figure 2.

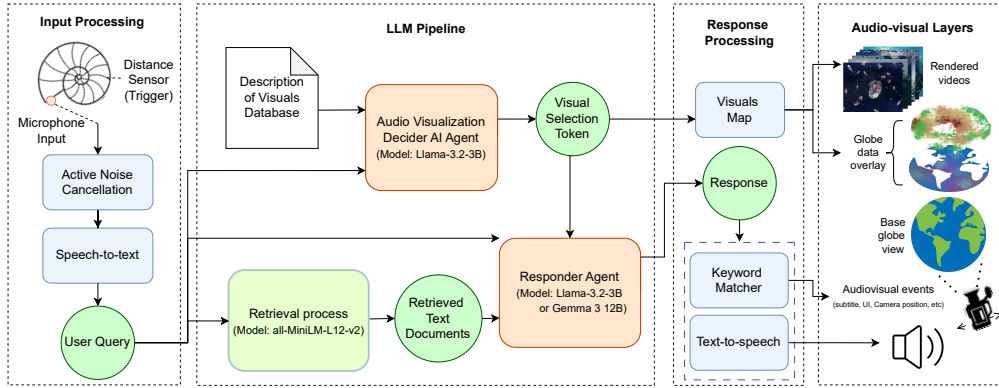


Figure 2: Overall diagram of the Sensorium Arc system

The overall structure of the proposed system is modularized into Input Processing, LLM Pipeline, Response Processing, and Audio-visual Layers.

- **Input Processing:** We designed and built an interaction module inspired by the Nautilus motif, implemented with Arduino-based control and incorporating a distance sensor, LED illumination, and a motor housed within a Nautilus shell. The Nautilus—often described as a living fossil bridging ancient and modern seas—serves as the focal point of interaction. When the participant approaches the interface, the distance sensor detects whether the user is within 50 cm. As the participant draws near, blue light ripples across its shell, signaling readiness to engage. If the proximity threshold is met, the microphone system is activated with active noise cancellation to suppress ambient sounds, and begins recording the participant’s voice. Fig. 1-b,c shows the default status and activation of the Nautilus interface. Recording continues until the user moves beyond the 50 cm range. The captured audio is then processed for speech-to-text conversion (4) and passed into the LLM pipeline for query interpretation and response generation.

• **LLM Pipeline:** LLMs appeared well-suited for Sensorium’s goal of integrating retrieval, multimodal control, and persona-driven narrative into a single conversational framework; however, early trials revealed that a single model handling all tasks offered reduced fine control, and suffered from issues such as prompt interference and opaque errors. Recent surveys on LLM-based multi-agent systems highlight that decomposing complex workflows into specialized, role-driven agents can improve performance, debuggability, and adaptability compared to monolithic designs (31). In such systems, each agent operates with tailored prompts, context management, and communication protocols, enabling optimization for its specific subtask while reducing prompt interference and role confusion.

This modularity supports Sensorium Arc’s need for combining retrieval, multimodal control, and persona-driven narrative in a coordinated yet flexible architecture. Each of these agents are stateless and lightweight where possible, with more intensive reasoning and complex context maintained by the final response generation step. The implementation builds upon and extensively modifies the open-source LLMUnity framework (7). Our multi-agent architecture includes the Visualization Decider Agent, Retrieval with Query Rewriter Agent, and Responder Agent.

Visualization Decider Agent - The prepared globe datasets, gathered from NASA EarthData (19), and Unity-rendered video layers become central to the experience when active. To ensure that the Ocean’s responses can adapt to these visuals, we moved the selection process upstream in the LLM pipeline. Specifically, we implemented a small stateless decision agent prior to the main response step, responsible for outputting a visual selection token based on the user query and a description of each prepared visual. For this agent, few-shot examples embedded in the system prompt, combined with explicit reasoning before the final output, improved performance and reliability with the Llama 3.2 3B model (5). We further applied GBNF grammar constraints to limit the output format and ensure that the final selection token matches only the available options.

Retrieval and Query Rewriter Agent - We adopted a flexible, simple RAG approach for both factual retrieval concerning our prepared data visualizations, as well as to shape and ground the Ocean persona in a curated corpus of textual resources, featuring essays, manifestos, transcripts, and other materials from the Harrisons’ archive (23). Among the documents are *The Time of the Force Majeure* (21), a volume surveying their work over several decades; “Apologia Mediterraneo,” a poetic dialogue with the Mediterranean created as part of *Counter Extinction V*; and a detailed catalog for *Peninsula Europe I*, which reflects on the consequences of climate change for Europe’s drainage basins. The unstructured resources were pre-processed into paragraphs, further split into sentences, embedded into a 384-dimensional vector space using the local embedding model all-MiniLM-L12-v2 (6), and indexed in a lightweight local database for retrieval.

In the retrieval step, queries were embedded with the same model and compared against the vector database to identify the k most similar sentence embeddings, and return both the matched sentences and their source paragraphs. Though we hope to explore more sophisticated methods of retrieval in the future, we used a low value of $k = 2$ after initial testing revealed that sometimes retrieving more paragraphs often began to fill up the final LLM step’s context window and muddle the response generation. This search used Approximate Nearest Neighbors (ANN) using the usearch library (32), via the LLMUnity plugin (7).

An additional LLM agent was incorporated into the retrieval process prior to semantic search, tasked with reformulating user queries to better align with the textual resources and removing unnecessary pragmatic markers. This design adapts the query rewriter proposed by Ma et al. (26). The agent is primed with general knowledge of the Harrisons’ texts and ecological philosophy contained in the database, enabling it to distill the conceptual core of philosophical queries and produce precise, document-relevant rewrites.

Through comparative testing on a set of representative queries, we selected Qwen 8B (8) for its strong rule-following behavior. To improve output quality, we applied basic chain-of-thought prompting (35), allowing the model to reason about the query before producing the final rewrite.

Responder Agent - As the final step in the LLM pipeline, this agent takes as context a description of the selected central visual (if any), two paragraphs from the retrieval process, and the user query. It outputs the Ocean’s final response. Tuning the system prompt for this agent proved most difficult, as it needed to establish our complex requirements for the Ocean persona in connection to the retrieval process mentioned above, including a primer on the key ideas articulated across the textual resources. It also needed to introduce the context that would be provided with each invocation and how it should

be used in the response generation process. We found that a well-formed zero-shot prompt worked well in most cases, but did not guarantee consistent response characteristics such as length or format.

- **Response Processing:** From there, response processing systems map keywords and tokens in the output to control additional audio-visual layers. The visual selection token is used to index into a map of all prepared data visualizations and video overlays. The Ocean response text output is passed through a simple modular keyword matcher, allowing addition of keyword event triggers for different audio-visual layers. The final response is also converted to an audio clip in this stage using the Jets text-to-speech model in Unity Sentis (3).

- **Audio-visual Layers:** Audio-visual layers are directly tied to the user’s interactive experience, combining data visualization, event triggers, video playback, and text-to-speech (TTS) audio accompanied by synchronized subtitles for accessibility. Scientific datasets, including atmospheric CO₂ emissions, chlorophyll concentration, sea surface temperature, ocean currents, and hyperspectral diffuse attenuation coefficients (Kd) from NASA Earthdata (19) and earthaccess (18), were preprocessed into globe-ready formats optimized for visual clarity, even when multiple datasets overlap. Keyword-based events trigger globe visualizations tied to specific locations and times, prompting camera movements or activating relevant visual layers. All videos regarding plankton blooms, ocean acidification, plastic waste dispersion, and sea-level rise were pre-rendered in Unity for performance efficiency, given the constraints of the available computing resources. While the current system relies on video playback, future iterations aim to replace videos with real-time renders, enabling user interaction with narratively-driven scenes.

By integrating these components seamlessly, the system allows participants to experience their whispered queries reaching the Ocean, which responds not only in voice but also with dynamically relevant audiovisual output (Figure 3). In designing these outputs, we considered both scientific accuracy and symbolic resonance.




Figure 3: Screenshots of the different exemplary scenarios (a) Default globe view before user interaction. (b) Chlorophyll concentration data visualized on the globe surface. (c) Water clarity rendered using 16 stacked layers of hyperspectral diffuse attenuation coefficient (Kd) data from NASA’s PACE satellite (2024). (d) Combined visualization of atmospheric CO₂ levels and ocean surface wind flow, highlighting multimodal environmental interactions.

- **Hardware and Runtime Environment:** Development and testing were conducted on Unity version 6000.0.24f1 running on Windows 11 with an Intel Core i9-13900HX CPU, 64 GB RAM, and an NVIDIA GeForce RTX 4090 Laptop GPU. CUDA acceleration was used to offload a configurable number of GPU layers from the CPU to improve LLM inference performance, resulting in average response latencies under 4 seconds after user input. Additional testing on macOS employed Vulkan for GPU acceleration. This local hardware configuration was chosen to ensure the system can be deployed in a mobile form factor suitable for public exhibition. The software is installable on both Windows and macOS to maximize accessibility, and we are planning educational installations for schools and public institutions.

4 Discussion

Human–AI interaction in Sensorium Arc raises questions about the evolving role of humanity in technologically mediated ecological discourse. By engaging participants in an embodied dialogue with a machine that speaks as the Ocean, the work foregrounds distinct human capacities such as empathy, moral reasoning, and cultural memory. These qualities are placed in dialogue with the machine’s strengths in data retrieval, multimodal synthesis, and maintaining a consistent ecological persona. This interplay reframes sustainability as not merely a technical or environmental metric, but as a shared cultural practice shaped through co-authorship with non-human agents. In doing so, it invites reflection on how human values, ethical priorities, and creative agency are encoded into

AI systems, and how such systems might, in turn, influence collective ecological imagination and responsibility.

While the current implementation of Sensorium Arc successfully delivers an immersive ocean conversational experience, several limitations highlight opportunities for refinement. We identify four main directions: Improving RAG reliability, Optimizing LLM agent performance, Artistic advancement, and Expanding scientific and social impact.

Improving RAG reliability - Our existing query-rewriting agent (24) could be extended with query expansion, generating multiple semantically distinct rewrites per user query to increase retrieval diversity. Categorizing embedded documents by type (e.g., artistic media, scientific studies, activist narratives, data) may better guide the responder agent toward responses that are both poetic and scientifically grounded. To mitigate outdated information, an autonomous update process could embed new documents as they become available (15). We may also explore “Read the Doc before Rewriting” (R&R) approaches (34) if current rewriting proves insufficient.

Optimizing LLM agent performance - The need to integrate multiple context sources and maintain conversation history required a large model, which improved quality over smaller models such as Llama 3.2 3B but increased latency (5). Since LLM inference competes with Unity’s real-time rendering for GPU resources, future deployments may offload inference to the cloud to alleviate hardware bottlenecks and support a wider range of devices. Moreover, recent study (16) highlights open challenges in multi-agent LLM design, such as ensuring clear role definitions, maintaining inter-agent alignment, and verifying outputs. Addressing these challenges presents an opportunity for future iterations of Sensorium Arc to refine its architecture, improve coordination efficiency, and enhance robustness, advancing toward more scalable and reliable multimodal interaction.

Artistic advancement - Sensorium’s core reflects Newton Harrison’s belief that the World Ocean could be given a voice—an artistic as well as technical pursuit requiring deliberate direction. Using RAG for poetic grounding is one step, with further potential in narrative visualization, interactivity, and immersion. Future sequences might open on a vast ocean horizon, draw participants in with layered marine soundscapes, and transition from macro to micro views, such as an explorable coral reef, deepening engagement with the Ocean’s response.

Expanding scientific and social impact - Planned steps include controlled user studies on retention, empathy, and behavioral change; connecting emotional resonance with civic engagement; and integration of more diverse and multilingual marine datasets via open API frameworks.

5 Conclusion

Sensorium Arc integrates conversational AI, RAG, and immersive media to render ocean data as an affective, dialogic experience. By combining scientific accuracy with poetic narration, it invites participants into co-authored exchanges with a responsive Ocean persona. Exhibitions show that this embodied interaction can shift environmental data from static representation to participatory narrative. The modular multi-agent design supports integration of new datasets, languages, and contexts, offering a generalizable framework for human–machine–ecosystem interaction.

Beyond technical contributions, Sensorium Arc models human–AI collaboration where humans bring curiosity, empathy, and cultural storytelling, and AI offers access to scientific archives, multimodal synthesis, and a consistent ecological persona. This interplay reframes sustainability as a cultural practice co-shaped with non-human agents, embedding values into AI systems that speak for the more-than-human world and strengthening collective ecological imagination.

Acknowledgments and Disclosure of Funding

This work was supported by Kuali Research Award (314284-00001), the Arts for All, and the Immersive Media Design program at the University of Maryland, College Park. We also gratefully acknowledge the AlloSphere and JoAnn Kuchera-Morin for their insightful guidance and to the Getty Foundation supported initiative PST2024: Art/Science Collide for their support and encouragement. Original funding for Sensorium came from the Metabolic Studio of the Annenberg Foundation and the Joan and Irwin Jacobs Family Fund.

References

- [1] The AlloSphere Research Facility. <https://allosphere.ucsb.edu>, Accessed October 27, 2025.
- [2] Cool Science. <https://www.coolscience.net>, Accessed October 27, 2025.
- [3] inference-engine-jets-text-to-speech · Hugging Face. <https://huggingface.co/unity/inference-engine-jets-text-to-speech>, Accessed October 27, 2025.
- [4] inference-engine-whisper-tiny · Hugging Face. <https://huggingface.co/unity/inference-engine-whisper-tiny>, Accessed October 27, 2025.
- [5] meta-llama/Llama-3.2-3B · Hugging Face, Dec. 2024. <https://huggingface.co/meta-llama/Llama-3.2-3B>, Accessed October 27, 2025.
- [6] sentence-transformers/all-MiniLM-L12-v2 · Hugging Face, Jan. 2024. <https://huggingface.co/sentence-transformers/all-MiniLM-L12-v2>, Accessed October 27, 2025.
- [7] LLMUnity, Aug. 2025. <https://github.com/undreamai/>, Accessed October 27, 2025.
- [8] Qwen/Qwen3-8B · Hugging Face, Aug. 2025. <https://huggingface.co/Qwen/Qwen3-8B>, Accessed October 27, 2025.
- [9] Sea Level Rise Explorer Projects: Turner Station (MD), Santa Cruz (CA), Long Beach (CA), West Palm Beach (FL), and Stinson Beach (CA), 2025. <https://www.virtualplanet.tech/#projects>, Accessed October 27, 2025.
- [10] [Video] Sensorium Arc, 2025. <https://www.myunginlee.com/arc>, Accessed October 27, 2025.
- [11] S. J. G. Ahn, J. N. Bailenson, and D. Park. Short- and long-term effects of embodied experiences in immersive virtual environments on environmental locus of control and behavior. *Computers in Human Behavior*, 39:235–245, Oct. 2014. doi: 10.1016/j.chb.2014.07.025
- [12] M. Arslan, H. Ghanem, S. Munawar, and C. Cruz. A Survey on RAG with LLMs. *Procedia Computer Science*, 246:3781–3790, 2024. doi: 10.1016/j.procs.2024.09.178
- [13] J. O. Bailey, J. N. Bailenson, J. Flora, K. C. Armel, D. Voelker, and B. Reeves. The Impact of Vivid Messages on Reducing Energy Consumption Related to Hot Water Use. *Environment and Behavior*, 47(5):570–592, June 2015. doi: 10.1177/0013916514551604
- [14] J. Calil, G. Fauville, A. Queiroz, K. Leo, A. Mann, T. Wise-West, P. Salvatore, and J. Bailenson. Using Virtual Reality in Sea Level Rise Planning and Community Engagement—An Overview. *Water*, 13(9):1142, Apr. 2021. doi: 10.3390/w13091142
- [15] D. N. Campo, J. Conde, Alonso, G. Huecas, J. Salvachúa, and P. Reviriego. Real-time Spatial Retrieval Augmented Generation for Urban Environments, 2025. doi: 10.48550/ARXIV.2505.02271
- [16] M. Cemri, M. Z. Pan, S. Yang, L. A. Agrawal, B. Chopra, R. Tiwari, K. Keutzer, A. Parameswaran, D. Klein, K. Ramchandran, M. Zaharia, J. E. Gonzalez, and I. Stoica. Why Do Multi-Agent LLM Systems Fail?, Apr. 2025. arXiv:2503.13657. doi: 10.48550/arXiv.2503.13657
- [17] Y. Chen, W. Wang, S. Lobry, and C. Kurtz. An LLM Agent for Automatic Geospatial Data Analysis, Oct. 2024. arXiv:2410.18792. doi: 10.48550/arXiv.2410.18792
- [18] N. Earth Science Data Systems. NASA Earthaccess, Apr. 2025. <https://www.earthdata.nasa.gov/data/tools/earthaccess>, Accessed October 27, 2025.
- [19] N. Earth Science Data Systems. Your Gateway to NASA Earth Observation Data | NASA Earthdata, Aug. 2025. <https://www.earthdata.nasa.gov/>, Accessed October 27, 2025.
- [20] P. Green. Newton Harrison, a Founder of the Eco-Art Movement, Dies at 89. *The New York Times*, Sept. 2022.

- [21] H. Harrison, A. Douglas, N. Harrison, C. Fremantle, W. Fox, E. Heartney, R. Malina, P. Mankiewicz, D. Sagan, P. Kruse, and others. *The Time of the Force Majeure: After 45 Years, Counterforce is on the Horizon*. Prestel, 2016.
- [22] J. Kuchera-Morin. Parsing the Senses to Explore Complex Systems: Multimodal Investigation at the AlloSphere Research Facility. In *Integrative Contemporary Art and Science Practices*. Routledge, 2025.
- [23] M. Lee. Sensorium Conversational Corpus for Retrieval-Augmented Generation, Aug. 2025. doi: 10.5281/ZENODO.16777228
- [24] J. A. d. O. Lima. Improving RAG Retrieval via Propositional Content Extraction: a Speech Act Theory Approach, Mar. 2025. arXiv:2503.10654. doi: 10.48550/arXiv.2503.10654
- [25] X. Liu, Y.-W. Tai, and C.-K. Tang. ChatCam: Empowering Camera Control through Conversational AI. *Advances in Neural Information Processing Systems*, 37:54483–54506, Dec. 2024.
- [26] X. Ma, Y. Gong, P. He, H. Zhao, and N. Duan. Query Rewriting for Retrieval-Augmented Large Language Models, Oct. 2023. arXiv:2305.14283. doi: 10.48550/arXiv.2305.14283
- [27] V. V. Manivannan, Y. Jafari, S. Eranky, S. Ho, R. Yu, D. Watson-Parris, Y. Ma, L. Bergen, and T. Berg-Kirkpatrick. ClimaQA: An Automated Evaluation Framework for Climate Question Answering Models. 2024. doi: 10.48550/ARXIV.2410.16701
- [28] D. M. Markowitz and J. N. Bailenson. Virtual reality and the psychology of climate change. *Current Opinion in Psychology*, 42:60–65, Dec. 2021. doi: 10.1016/j.copsyc.2021.03.009
- [29] D. M. Markowitz, R. Laha, B. P. Perone, R. D. Pea, and J. N. Bailenson. Immersive Virtual Reality Field Trips Facilitate Learning About Climate Change. *Frontiers in Psychology*, 9:2364, Nov. 2018. doi: 10.3389/fpsyg.2018.02364
- [30] H. Rahimi, J. Cattoni, M. Beghili, M. Abrini, M. Khoramshahi, M. Pino, and M. Chetouani. Reasoning LLMs for User-Aware Multimodal Conversational Agents, 2025. doi: 10.48550/ARXIV.2504.01700
- [31] W. Shen, C. Li, H. Chen, M. Yan, X. Quan, H. Chen, J. Zhang, and F. Huang. Small LLMs Are Weak Tool Learners: A Multi-LLM Agent, 2024. doi: 10.48550/ARXIV.2401.07324
- [32] A. Vardanian, I. Nazaryan, Semantic Release Bot, C. Spence, and E. L. Gonçalves. unumcloud/usearch: v0.5.0, May 2023. doi: 10.5281/ZENODO.7949416
- [33] R. Wallen. Helen and Newton Harrison: Visionary EcoArtists. *Leonardo*, 56(3):312–313, June 2023. doi: 10.1162/leon_a_02386
- [34] Q. Wang, Y. Cao, Y. Liu, J. Zhao, and P. Luo. Read the Docs Before Rewriting: Equip Rewriter with Domain Knowledge via Continual Pre-training, 2025. doi: 10.48550/ARXIV.2507.00477
- [35] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models, Jan. 2023. arXiv:2201.11903. doi: 10.48550/arXiv.2201.11903
- [36] N. Wiratunga, R. Abeyratne, L. Jayawardena, K. Martin, S. Massie, I. Nkisi-Orji, R. Weerasinghe, A. Liret, and B. Fleisch. CBR-RAG: Case-Based Reasoning for Retrieval Augmented Generation in LLMs for Legal Question Answering. In J. A. Recio-Garcia, M. G. Orozco-del Castillo, and D. Bridge, eds., *Case-Based Reasoning Research and Development*, pp. 445–460. Springer Nature Switzerland, Cham, 2024. doi: 10.1007/978-3-031-63646-2_29
- [37] O. N. Yalcin, N. Abukhodair, and S. DiPaola. Empathic AI Painter: A Computational Creativity System with Embodied Conversational Interaction. 2020. doi: 10.48550/ARXIV.2005.14223