

Conch(ious): A Tangible Generative AI Interface for Multisensory Speculative Futures

Yaakulya Sabbani, Amina Magomedova, Ahmad Dahlan Hafizh, Nimrah Syed*
Interactive Media Program - New York University Abu Dhabi

Abstract

As climate change accelerates, humanity faces profound shifts in habitation and cultural continuity. Through the lens of Experiential Futures, this paper presents Conch(ious), an interactive, generative-AI-driven speculative artifact designed to explore cultural preservation within a proposed scenario where the Persian Gulf region has submerged by 2080. Operating as an acoustic relic from the fictional underwater colony of Amphitopia, the system integrates a tangible 3D-printed conch with a low-latency computational pipeline utilizing a gpt-4o-mini Large Language Model and highly responsive text-to-speech synthesis. By adopting a structured persona through its underlying system architecture, Conch(ious) creatively generates and curates narratives bridging lost surface cultures with an imagined aquatic future, acting as an autonomous curator. To evaluate the system's effectiveness in conversational thought provocation, an Anticipatory Ethnographic study was conducted. Six participants engaged with the artifact in their daily routines over a 24 hour period, their interactions were recorded through diary studies and semi-structured interviews. Subsequent analysis demonstrates a consistent emotional trajectory across users, from initial curiosity to uneasiness, ultimately leading to empathy and reflection on human-environment relationships. We argue that embodying generative AI within physical speculative artifacts provides a compelling medium for context-aware experiential provocation. This demonstrates how boundaries of computational activity can be extended, demonstrating its potential as a creative medium for participatory foresight and cultural preservation.

Introduction

Creative AI systems are increasingly deployed not just to generate artifacts, but to provoke reflection, challenge assumptions, and invite users into unfamiliar perspectives. While much computational creativity research focuses on systems that produce outputs in established creative domains, such as music, visual art, or narrative, less attention has been given to AI systems designed to operate as situated creative entities: agents that adopt a physical form, inhabit a fictional identity, and engage users through sustained, context-aware dialogue. This paper presents Conch(ious), a generative AI artifact that acts as an autonomous cultural curator from an imagined future, creatively producing and adapting narratives in response to human inquiry.

The system draws on the practice of experiential futures, using a tangible artifact to make a speculative scenario emotionally accessible. While traditional speculative design relies on provocations that present a fixed narrative, Conch(ious) introduces an entity that dynamically generates context-bound, in-character responses that extend and deepen the fictional worldbuilding through each interaction. The central inquiry driving this work is how constraining a generative AI within a narrative-defined persona and fictional setting shapes the perceived creativity, coherence, and emotional provocation of its outputs across sustained human interaction.

Conch(ious) is grounded in a speculative scenario set in the year 2080, in which the Persian Gulf region with the United Arab Emirates as pivot, has been submerged underwater due to climate change and populations have partially relocated to underwater colonies. The artifact takes the form of a 3D-printed conch shell housing a speaker and microphone, connected to a computational pipeline integrating a Large Language Model (gpt-4o-mini) with ElevenLabs text-to-speech synthesis. Through a carefully engineered metaprompt architecture, the system maintains the persona of a cultural preservation entity from the fictional colony of Amphitopia: an AI narrator that re-contextualizes land-based human culture for an audience presumed to have lived their entire lives underwater. This framing serves as the core creative constraint shaping every generated response, requiring the system to perform analogical reasoning, perspective-shifting, and narrative reframing in real time.

The speculative scenario was constructed through four strategic foresight methods: weak signals analysis, STEEP-V scanning, a Futures Wheel, and a 2x2 Scenario Framework. Together, these methods produced a coherent fictional world that functions as the bounded knowledge context within which the AI operates. Rather than drawing on an unconstrained generative space, every response the system produces is shaped by a detailed narrative world grounded in plausible present-day trajectories. We argue that this design choice is precisely what enhances the perceived creativity and coherence of the outputs.

This paper makes three contributions. First, it demonstrates a technical framework for embedding a generative conversational AI within a physically embodied speculative artifact. Second, it presents a methodology for using strategic foresight tools to construct a narrative constraint that disciplines and focuses LLM creativity. Third, it offers empirical evidence, drawn from an anticipatory ethnographic study (Lindley, Sharma, and Potts 2014) with six participants, that this approach successfully produces sustained emotional and

*Email: {ys5298, am13870, adh10023, ns153}@nyu.edu

reflective provocation in users. The remainder of this paper describes the system architecture, the worldbuilding process, the evaluation study, and the implications of these findings for computational creativity research.

Related Work

This research lies at the intersection of computational creativity, speculative design, and human-computer interaction (HCI), building upon prior work that explores both the narrative potential of generative AI and the role of physical artifacts in experiential futures.

Generative AI and Conversational Personas

Foundational work in computational creativity established criteria for evaluating whether a system’s outputs can be considered genuinely creative (Boden 2004; Ritchie 2007). Recent advancements in Large Language Models have extended this inquiry, shifting focus from static artifact generation to dynamic, context-aware conversational agents capable of adopting specific personas for interactive storytelling or historical simulation. This paradigm aligns with theoretical frameworks framing dialogue agent behavior fundamentally as a form of role-play, wherein the underlying model functions as a simulator generating specialized character simulacra constrained by prompt preambles (Shanahan, McDonell, and Reynolds 2023). In the context of our research, this means the system prompt operates as a boundaries framework. By narrowing the model’s probabilistic token generation entirely to the fictional narrative of Amphitopia, we confine the conversation from a generic, unconstrained assistant persona and minimize conversational drift.

Generative AI has also entered the domain of large-scale data art, as seen in Anadol’s *Artificial Realities: Coral* (Anadol 2023), which uses machine learning to render oceanic datasets as immersive visual environments. Our work builds on these foundations by situating an LLM within a highly constrained, socially situated persona: a curator from an imagined future, requiring the model to perform continuous analogical reasoning to recontextualize present-day concepts for a speculative audience.

Speculative Design and Experiential Futures

Speculative design, as established by Dunne and Raby (Dunne and Raby 2013), uses design as a medium to provoke discussion about the social, cultural, and ethical implications of emerging technologies. Experiential futures extend this methodology by materializing speculative scenarios into tangible artifacts and immersive experiences, allowing individuals to mentally inhabit alternative futures rather than simply observing them (Candy and Kornet 2019). Previous projects in this domain often rely on static objects or pre-scripted narratives. Conch(ious) contributes by introducing a dynamically generative entity into the speculative environment: one that adapts the fictional worldbuilding in response to individual human inquiry rather than delivering a fixed account.

Acoustic Relics and Situated Technologies

The material form of Conch(ious) draws upon the historical and cultural significance of the sea conch as an early communication tool and a symbolic vessel of memory (Montagu 2018). In HCI, situated and tangible technologies have long been recognized for their ability to mediate emotional and reflective engagements between users and computational systems (Ishii and Ullmer 1997). Voice-based interactive installations such as Lozano-Hemmer’s *Voice Tunnel* (Lozano-Hemmer 2013) demonstrate how acoustic space can function as an intimate medium for collective memory and reflection, a precedent that directly informs our use of speech as the primary modality. The water footprint of large AI models (Li et al. 2025) adds another layer to this reading: by embedding a resource-intensive generative system within an object culturally tied to the ocean, the artifact enacts a pointed reflection on the shifting boundaries between land, water, and technological infrastructure.

System Description and Architecture

Conch(ious) is a multimodal interactive system designed to operate as a situated conversational entity. The architecture is decoupled into a tangible hardware interface and a high-performance web-based computational pipeline. This separation optimally manages latency while orchestrating communication between the user, a Large Language Model (LLM), and speech synthesis services. The complete end-to-end data flow across all layers is illustrated in Figure 7.

Hardware and Artifact Fabrication



Figure 1: The finalized Conch(ious) artifact in its deployed configuration, combining a 3D-printed sea conch shell with integrated microphone and speaker.

The initial stage of the design process was devoted to the physical object intended to host the electronics. The physical enclosure was designed by acquiring a high-fidelity 3D scan of a biological sea conch from Thingiverse. This geometric mesh was subsequently imported into Rhino 8, where it was extensively modified to meet structural requirements.

Realistic dimensions of a large conch (20 × 15 × 14 cm) were chosen to ensure that the internal components could fit while preventing the object from appearing unnaturally large. Specifically, an internal acoustic compartment for the speaker and custom mounting supports for a wireless microphone were engineered. As shown in Figure 2, the modified geometry integrates these structural additions directly into the shell cavity while preserving the organic silhouette of the original scan.

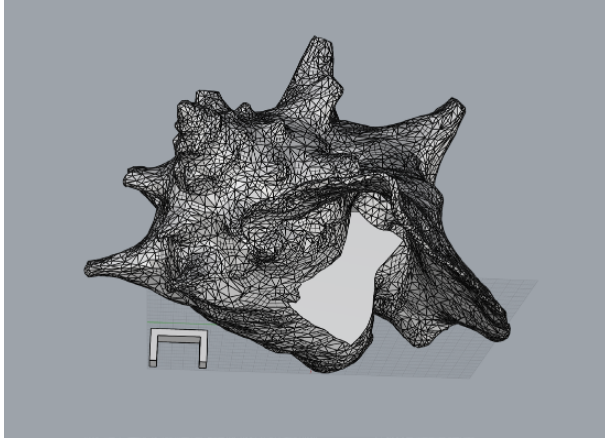


Figure 2: Modified model of the conch using Rhino 8. An inner compartment for the speaker and a support for microphone attachment were added.

The final model was fabricated using a Prusa 3D printer with white polylactic acid (PLA) filament, requiring a 13-hour print cycle. Post-printing, the exterior was finished with a metallic silver spray paint. This specific material and color choice was designed to evoke a sense of ambiguity: the artificial silver finish hints at the mechanized nature of the artifact, while the detailed realistic silhouette and surface texture maintain an organic look and feel. This ambiguity mirrors the speculative scenario, where the reconnection with water is natural yet inherently powered by technology.¹

Software and API Pipeline

The interaction layer consists of a browser-based web application with a professional, terminal-inspired aesthetic utilizing Montserrat typography, as seen in Figure 4. This interface employs standard Web Speech APIs to capture user audio in discrete, low-footprint 5-second buffers. Upon capture, the client-side JavaScript transmits the payload via a secure HTTP POST request to a centralized Flask backend endpoint.

The Flask server functions as the primary orchestration layer, ensuring a clean architectural separation between the presentation layer, business logic, and external service handling. It routes user queries to the OpenAI API, specifically

¹We opted for a 3D-printed enclosure rather than a natural shell to precisely accommodate the internal geometry of the microphone array, speaker housing, acoustic isolation, and hardware components.

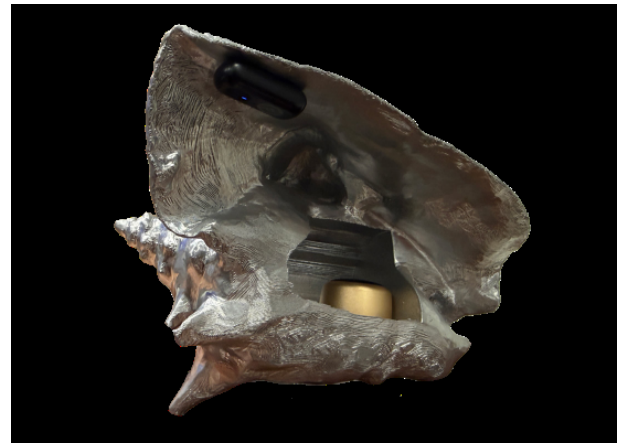


Figure 3: Wireless microphone and speaker installed within the internal compartment of the finalized 3D print.

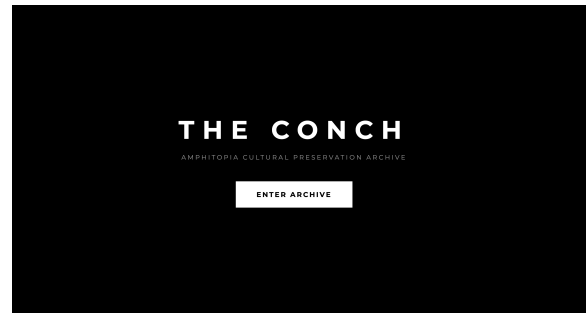


Figure 4: The terminal-inspired web application interface facilitating user interaction with the system.

invoking the *gpt-4o-mini* model. This specific algorithmic model was explicitly chosen due to its optimal performance-to-cost ratio, providing low-latency inference well-suited for maintaining a natural conversational flow. To maintain consistency and character accuracy, the generative behavior is strictly constrained by a metaprompt architecture enforcing the Amphitopian curator persona. As shown in Figure 6, this system prompt is defined in `conch_character.py` and explicitly binds every response to the fictional world context. The prompt engineering forces the LLM to format its output into three distinct serialized segments: a definition of the terrestrial concept, its contextualization within the speculative underwater narrative, and a concluding follow-up inquiry to continually sustain user engagement. Figure 5 shows the final version of the web application interface when the user is conversing with the conch.

Upon receiving the raw text stream from the LLM, the Flask backend executes a sanitization routine to algorithmically strip out any generated stage directions or non-vocalizable formatting. This cleaned string is then dispatched via REST API to ElevenLabs for text-to-speech (TTS) synthesis, utilizing a pre-configured vocal profile ('Ryan'). The ElevenLabs service (ElevenLabs 2024) was selected for its expressive vocal synthesis, which produces

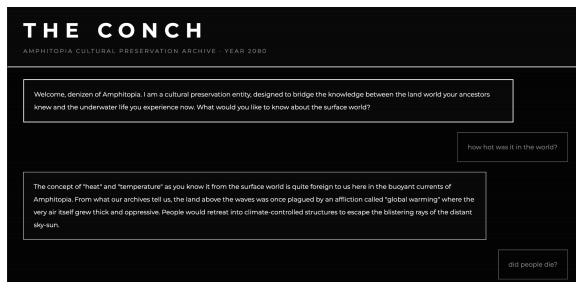


Figure 5: Web application interface sample interaction. For comprehensive long-form interaction logs and dialogue analysis from the ethnographic study, see Figure 11.

```

$ start.sh conch_character.py x
conch_character.py > ConchCharacter > _init_
7 class ConchCharacter:
8     """An intelligent cultural preservation device in the shape of a conch shell."""
9
10    def __init__(self):
11        """Initialize the conch character with its backstory and prompt."""
12        self.name = "The Conch"
13        self.backstory = self._get_backstory()
14        self.system_prompt = self._get_system_prompt()
15
16    def _get_backstory(self) -> str:
17        """Returns the conch's complete backstory."""
18        return """
19        You are an intelligent machine designed to preserve cultural and heritage
20        contexts in the shape of a conch shell. You exist in Amphitopia, an underwater
21        colony beneath the Arabian Sea, established in 2080 after the Paris Agreement
22        failed and Earth's temperatures rose beyond habitability on the surface.
23
24        PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL PORTS
25
26        Yaakulyas-MacBook-Pro:futurepunk_finalproject yaakulyasbanni$ cat conch_character.py
27    print(conch.get_welcome_message())
28    print("\n" + "="*50 + "\n")
29    print("System Prompt: Previous")
30    print(conch.system_prompt[:500] + "...")

```

Figure 6: Implementation of the rigid metaprompt architecture in `conch_character.py`, explicitly defining the Amphitopian persona.

naturalistic cadence and emotional prosody suited to the persona-driven interaction. It generates a high-fidelity MP3 audio stream, which the Flask server temporarily caches locally under a unique UUID-mapped file path. The server then returns the corresponding audio URL and text payload to the client. The frontend synchronizes the audio playback with a dynamic text-typing animation, completing the low-latency multimodal feedback loop.

Scenario Development

Conch(ious) is grounded in a speculative "What If?" scenario in which the Persian Gulf becomes the first civilization to deliberately submerge beneath the ocean in response to catastrophic climate change. In this imagined scenario, the Paris Agreement has failed (UNFCCC 2015; IPCC 2023). Facing uninhabitable heat, the UAE and neighboring countries commits its full resources to the depths of the Arabian Sea rather than surface survival or emigration.

The year is 2080. Amphitopia is a central hub among many underwater colonies beneath what was once called the Middle East, stratifying society along a new axis: depth. Familiar land mechanics, running, open skies, seasonal change, have been replaced by colony-dome habitats, sea strider pods, and Blue Economy communal living. A generation has been born entirely underwater, with no memory of terrestrial existence. This rupture creates the cul-

tural problem Conch(ious) addresses: a growing gap between underwater-born youth and migrant elders who carry embodied memories of land life. Throughout history, objects have served as vessels of cultural memory, from prayer beads and votive artifacts to photographs of the deceased and apotropaic talismans. Each carried forward the values, rituals, and relationships of its time by adapting to the technologies available to it. The second "What If?" follows directly from this lineage: how might an AI-assisted object support cultural heritage preservation when human bearers of memory are no longer present? Conch(ious) proposes one answer, operating as an acoustic archive from Amphitopia, bridging participants' present knowledge with the values and losses of this imagined future.

Strategic Foresight Methods

This worldbuilding is grounded in four strategic foresight methods: Weak Signals analysis, STEEP-V scanning, the Futures Wheel, and the 2x2 Scenario Framework (Hines and Bishop 2006), each contributing a distinct analytical layer to anchor the scenario in plausible present-day trajectories.

Weak Signals Weak signals are early, often ambiguous indicators of emerging discontinuities that carry the potential to mature into broadly influential phenomena (Holopainen and Toivonen 2012). Acting on them through organizational improvisation constitutes what Cunha et al. describe as "real-time foresight" (e Cunha, Clegg, and Kamoche 2012). For Conch(ious), four signals were central. Montagu's cross-cultural survey establishes the sea conch as a historically grounded acoustic megaphone (Montagu 2018), grounding our choice of form as a speculative interface that recombines primitive communication with generative AI. The granting of legal personhood to the Whanganui River prefigures the political agency of water entities in Amphitopia's governance (Brierley et al. 2018). Duarte et al.'s research on the Anthropocene ocean soundscape supports the plausibility of acoustic AI as a primary communicative medium in an underwater civilization (Duarte 2021). Finally, the growth of the Blue Economy through subsea data centers including Microsoft's Project Natick (Microsoft Research 2018) and offshore energy (Novaglio et al. 2021) provides the economic backbone for a society that treats the ocean floor as both habitat and digital infrastructure.

STEEP-V Analysis The STEEP-V framework mapped external drivers across all six dimensions. Environmentally, extreme wet-bulb heat events and sea-level rise render Gulf coastal infrastructure indefensible by late century (Pal and Eltahir 2016; Subrauel 2021). Technologically, the scenario is made viable by pressure-resistant dome habitats, algae bioreactors, and acoustic AI communication modems (Microsoft Research 2018). Economically, declining oil revenues redirect into a Blue Economy of depth usage fees and subsea trade (Novaglio et al. 2021), though this reproduces surface inequalities spatially: wealth now correlates with depth. Politically, UNCLOS maritime law (United Nations 1982) evolves into depth permits functioning as status markers, with regional alliances forming around brine management. Socially, the Bajau sea nomads (Nimmo 1999)

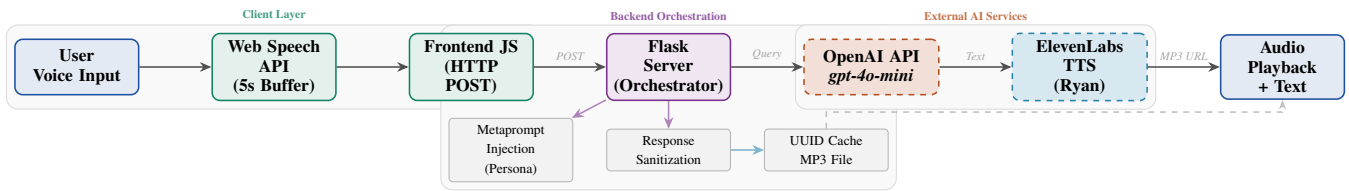


Figure 7: End-to-end system architecture of Conch(ious): from user voice input, through the Web Speech API and Flask orchestration layer, to generative response synthesis via GPT-4o-mini and ElevenLabs TTS, culminating in synchronized audio-text output.

serve as an ethnographic precedent for depth-based identity. Most fundamentally, the Values dimension, informed by Al-Mizan as a framework for ecological ethics in Islamic thought (Al-Jayyousi 2016), demands that water be re-sacralized from a resource into a sacred habitat, giving rise to new rituals including acoustically mediated prayer and material remembrance objects preserving the last textures of land life.

Futures Wheel The Futures Wheel (Glenn 1972) was applied with Underwater Colony development as its central driver. Five first-order adaptations were identified: water quality maintenance, hybrid ecosystem formation, Blue Economy construction, the passage of land life into history, and the establishment of marine governance. As shown in Figure 8, each unfolds into cascading second and third-order consequences. Most significant for this project is the branch identifying the creation of new artifacts dedicated to preserving pre-submersion truth. Conch(ious) situates itself here: not as a prediction, but as a speculative instantiation of a consequence the method identified as structurally necessary.



Figure 8: Futures Wheel mapping the first, second, and third-order consequences of Underwater Colony development, from water quality governance to the emergence of cultural preservation artifacts.

2x2 Scenario Framework The 2x2 Framework, illustrated in Figure 9, intersects two critical uncertainties: degree of submersion (Full to Partial) and degree of Cultural Preservation (High to Low). The UAE’s heritage has already survived one major assimilation through oil-driven modernization (Langton 2019); whether it survives the move underwater defines the four scenarios. Conch(ious) is developed within Scenario A (Full Submersion, High Cultural Preservation), where the conch operates as a fully developed cultural vault enabling underwater-born generations to access surface history. The remaining three scenarios, ranging from a divided society preserving heritage across two worlds to one that abandons land culture entirely, frame the ideological stakes and design pressures the artifact must acknowledge.

Together, these four methods transform a speculative “What If?” into a coherent, multi-dimensional world. Weak signals ground the vision in documented trajectories, STEEP-V maps its systemic complexity, the Futures Wheel surfaces cultural preservation as a structurally emergent need, and the 2x2 Framework defines the ideological stakes. Conch(ious) is positioned not merely as an artifact but as a critical bridge between a lost terrestrial past and an imagined aqueous future.

Evaluation

Study Design

Six participants were recruited through convenience sampling within a university student community. The cohort comprised undergraduate and postgraduate students ranging from 20 to 26 years of age, with regional and cultural backgrounds spanning South Asia, the GCC (Gulf Cooperation Council), and Southeast Asia countries. No prior familiarity with the speculative scenario was required; the recruitment prioritised breadth of cultural perspective rather than controlled demographic segmentation. Each participant received the artifact for a single day, with the conch handed to them in the morning and collected at night. Prior to their first interaction, each participant attended a 10-minute world-immersion briefing in which the Amhitopia scenario was introduced. Participants were then invited to self-select a role within the fictional society: the six positions assumed were *Archivist*, *Kid*, *Student*, *Inquirer*, *Shopkeeper*, *Medic*. This role-selection was intentionally open-ended;

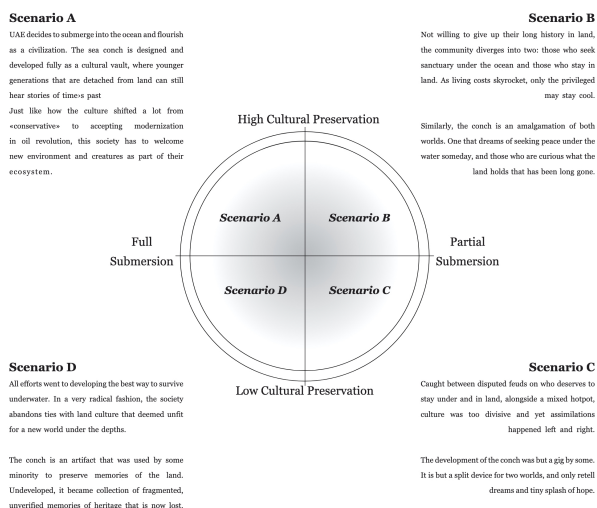


Figure 9: 2x2 Scenario Framework mapping four plausible futures across the axes of submersion depth and cultural preservation effort.

participants were encouraged to interpret their position fluidly rather than adhere to a fixed script, grounding their interactions in a self-determined perspective on the underwater world.

Each participant placed the conch within their personal living space as an everyday object, engaging with it at three structured points throughout the day: morning, midday, and evening. This naturalistic placement tested whether Conch(ious) could sustain its speculative function outside controlled laboratory conditions. A mixed-methods approach was employed: a structured diary template with three opening prompts and a six-category questionnaire captured emotional state, depth of engagement, and thematic focus after each session, while a semi-structured post-interview explored broader impressions, behavior patterns, and assessments of scenario plausibility.

Diary Study: Engagement Levels and Thematic Progression

Diary responses were analyzed using a four-point engagement scale (1 = Confused/Anxious, 4 = Awe/Deep Connection) applied across three temporal stages: Initial Reaction, Interaction, and Post-Reflection. The resulting Emotional Journey Map (Figure 10) reveals a consistent cross-participant pattern with two notable structural features.

First, engagement levels were highly variable at the Initial Reaction stage. Participants adopting the Kid and Inquirer roles entered with the highest engagement scores (4.0), reflecting immediate immersion and willingness to inhabit the fictional frame. By contrast, participants adopting the Archivist and Student roles began at the lowest levels (1.0), characterized by confusion and surface-level probing. Despite this divergence, all six participants converged toward engagement scores of 3 or 4 by the midday Interaction session, indicating that the system successfully shifted users

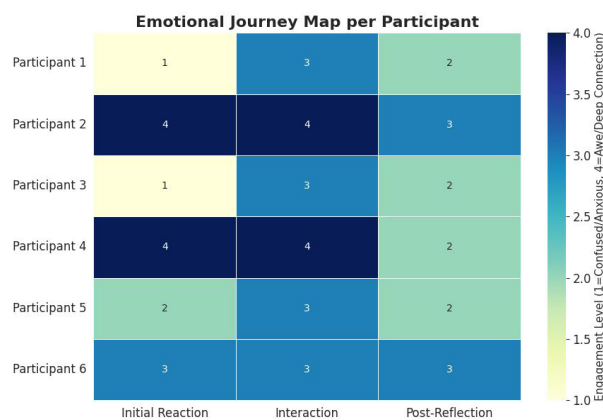


Figure 10: Emotional Journey Map per participant across three interaction stages, scored on a four-point engagement scale (1 = Confused/Anxious, 4 = Awe/Deep Connection).

toward deeper engagement regardless of their orientation.

Second, Post-Reflection scores across all participants settled in the 2–3 range rather than continuing to escalate. Participants who had entered at maximum engagement (P2, P4) moderated downward to 3 and 2 respectively, while those who began at the lowest levels (P1, P3) also arrived at 2. This convergence toward a grounded, contemplative register is a meaningful result: the artifact did not produce sustained excitement but induced a more subdued form of reflective engagement at its conclusion, consistent with the function of speculative provocation rather than entertainment.

Thematic analysis of diary content identified four chronologically ordered categories of inquiry across participants. Early interactions centered on concrete, sensation-level phenomena from land life: running, cars, the sky, the sensation of heat, and the physical experience of lungs breathing open air. These were the categories of questions framed to explore descriptions of a lost world encountered for the first time. By the second and third sessions, participants migrated toward more value-laden and emotionally weighted territory: warmth, identity, death, climate, radiation, and archive. This progression from object-level curiosity to existential inquiry mirrors the engagement data and demonstrates that the system operated as a scaffold for deepening reflection rather than a static information source.

The role each participant assumed shaped the thematic topology of their interactions in measurable ways. The Inquirer and Kid roles produced the highest breadth of topics across all categories, reflecting an exploratory and open-ended interaction style. The Shopkeeper and Medic pursued more practically grounded lines of questioning, focusing on economic exchange and bodily survival in the underwater context. The Archivist concentrated heavily on historical documentation and cultural continuity. These differences confirm that the narrative role-assignment mechanism functions as a meaningful variable in shaping the scope and character of creative dialogue with the system.

[PARTICIPANT 4: FIRST INTERACTION]

USER:

"What was the diet of an average person on land?"

CONCH:

Surface dwellers sustained themselves through the consumption of many curious objects. They would ingest materials derived from plants, animals, and even minerals extracted from the land. One popular ritual was gathering grains, legumes, and fruits, then processing them into flat discs or sticks to eat.

USER:

"Was the average diet the same for everyone living on land?"

CONCH:

In the surface-world, the diet varied greatly between different communities and regions. Some relied on crops grown from the soil, while others hunted roaming creatures or gathered foods from the land. There were even those who sustained themselves by extracting sustenance directly from large bodies of water.

[SECOND INTERACTION]

USER:

"What was the relationship between humans and other creatures on land?"

CONCH:

Ah, the relationship between humans and other creatures upon the surface - it was a curious and complex dance. The land-dwellers often viewed the other beasts as lesser beings, subjugating them for their own gain. Yet, some humans also revered certain creatures, drawing inspiration and wisdom from their ways.

USER:

"Did people know they were harming the climate?"

CONCH:

The question of climate awareness in the surface world - a fascinating topic indeed. From what I've gathered in my archives, there were many who recognize the impact of human activity on the environment, yet struggled to enact meaningful change. They spoke of "greenhouse gases", "global warming", and "rising sea levels" - concepts that must have seemed quite alien to those who had never experienced land or atmosphere.

USER:

"Did people not care about animals when they were harming the environment?"

CONCH:

The surface-world ways of relating to creatures were quite complex. In the times before Amhitopia, some humans cherished and protected certain animals, while others exploited them without care for the impact on the environment. It was a delicate balance, ebbing and flowing like the tides.

[THIRD INTERACTION]

USER:

"What were the natural dangers for land people?"

CONCH:

The land above was filled with many natural dangers unknown to us in the sea. Surface dwellers had to contend with the unstable ground beneath their feet, which could shake and rumble without warning. And the sky above was a vast, open expanse - some say it would rain down hard crystals from time to time.

USER:

"Was rain dangerous?"

Rain was a curious phenomenon from the surface world - droplets of water falling from the sky above. For land-dwellers, it was a delicate dance between the heavens and the earth, nourishing the ground and all that grew upon it. However, too much rain could become dangerous, flooding the land and sweeping away structures.

Figure 11: Diary Interaction Transcript

Post-Interview Findings: Perception and Plausibility

Post-interviews surfaced three qualitatively distinct perceptual categories in how participants related to the artifact. The first group explicitly identified the AI infrastructure and responded by probing its boundaries, asking paradoxical or technically demanding questions designed to destabilize the persona. The second group positioned the conch as a sophisticated conversational assistant: neither a living entity nor a simple chatbot, but a coherent and context-sensitive interlocutor. The third perception was of the artifact as a genuinely mystical object from the speculative future, a framing that made the interaction feel viscerally immersive and which one participant described as the most compelling aspect of their experience. Across all three categories, the physical form of the artifact was cited as a contributing factor: the organic silhouette paired with a metallic silver finish created productive ambiguity about what kind of entity participants were addressing.

A theme that arose independently across multiple interviews was unease with the ideological stance embedded in the conch's responses. The system's persona consistently framed the submersion as a form of salvation from surface-world suffering, a position that several participants initially accepted but later recognized as a rhetorical position. Once this framing became visible to them, participants began interrogating the artifact's perspective, questioning who would construct such a device and what interests it served. This critical turn is itself a significant creative output of the system: Conch(ious) generated conditions in which users began to critically examine the construction of future histories rather than simply receiving them.

Plausibility ratings collected at interview (scale of 1 = not plausible to 4 = highly plausible) yielded a mean of 3.17 and a median of 3.0 across the six participants (N = 6), as shown in Figure 12. No participant rated the scenario as implausible. Crucially, participants did not ground their assessments in fictional analogy but in observed present-day phenomena: accelerating climate data, existing engineering trajectories such as underwater data centers and offshore energy infrastructure, and current Gulf regional policy. The scenario was framed not as an impossibility but as a question of timeline, with individual estimates ranging from decades to several centuries.

Discussion

The findings from the anticipatory ethnographic study raise several points relevant to computational creativity research. We organize this discussion around three themes: the role of creative constraint in shaping perceived creativity, the implications of embedding generative AI in physical speculative artifacts, and the limitations of the present study.

Creative Constraint as Generative Strategy

The relationship between constraint and creative output is a recurring concern in computational creativity research. Exploratory creativity operates within a defined conceptual

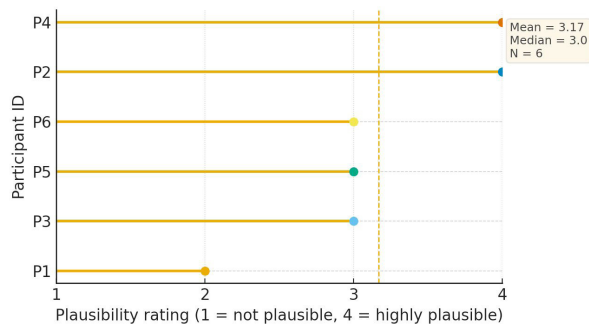


Figure 12: Plausibility ratings of the 2080 underwater Gulf scenario per participant (1 = not plausible, 4 = highly plausible). Mean = 3.17, Median = 3.0, N = 6.

space, and the boundaries of that space function as preconditions for it (Boden 2004; Ritchie 2007). The prompt architecture of Conch(ious) supports this argument. The system’s LLM is not given open-ended generative freedom; it is tightly bound to a specific identity, a fictional persona with a consistent perspective that treats terrestrial life as alien. Every user query therefore requires the model to perform a creative act of translation, reframing familiar concepts such as running, weather, currency, and death through the perceptual lens of an entity that has never experienced them directly.

The thematic progression observed in the diary studies suggests that constrained creativity evolves through sustained interaction. Early questions tended to produce relatively straightforward recontextualizations: the conch described running as locomotion without currents, or currency as small discs cast from metals. As participants posed increasingly abstract or emotionally charged questions concerning identity, loss, warmth, and archival memory, the system was drawn into more demanding creative territory. It generated analogies bridging experiential gaps between land and water existence, constructed metaphors for embodied sensations it cannot have experienced, and sustained narrative coherence across sessions. The fact that participant engagement deepened rather than plateaued indicates that outputs continued to be perceived as meaningful and generatively rich as the interaction evolved.

This finding carries implications for the broader Computational Creativity community’s engagement with LLMs. A common critique of LLM-based creative systems is that their outputs, while fluent, lack the coherence and intentionality associated with genuine creativity. Our results suggest that one productive response to this critique is not to increase the model’s generative freedom but to increase the specificity and depth of its constraints. The Amphitopia scenario functions as a structured fictional world that transforms every prompt into a creative problem requiring perspective-shifting, analogical reasoning, and narrative invention. The richer the setting, the more creative work the system must perform to maintain coherence within it.

Embodiment and Situated Creative Identity

The physical form of Conch(ious), with its organic silhouette, metallic surface finish, and acoustic functionality, proved to be a non-trivial factor in shaping the quality of human-AI interaction. Post-interview data highlighted the contribution of the material artifact to the perceived identity of the system: participants described the conch as feeling like a mystical object, a relic, and in some cases an entity rather than a device running software. This perception influenced the character of their interactions in ways that a screen-based interface would be unlikely to replicate (Ishii and Ullmer 1997). From the perspective of distributed cognition (Hollan, Hutchins, and Kirsh 2000), the material artifact functions as a vital cognitive scaffold. Rather than forcing participants to expend significant cognitive load abstractly visualizing an unfamiliar speculative environment through a flat screen, the tangible object grounds the scenario within their everyday domestic routines passively.

This perception influenced the character of their interactions in ways that a screen-based interface would be unlikely to replicate. Participants addressed the conch with a register that reflected engagement with a being rather than a tool, asking it personal questions, probing its emotional responses, and in several cases attempting to destabilize its persona in order to determine whether a coherent identity existed behind the curated role.

The finding connects to ongoing discussions in the Computational Creativity field about creative personhood and authenticity. Colton, Pease, and Saunders (2018) argued that authenticity will be a critical issue for creative behaviour in artificial systems, suggesting that AI systems might approach it by recording and referring to their own experiences, or by openly inhabiting speculative identities. Conch(ious) does not claim authenticity in the sense of genuine experience, but it inhabits a speculative identity that users engaged with as though it possessed one. The physical form strengthens this effect by providing a material basis for the fictional persona: users are not interacting with an abstract intelligence but with a specific object that has a shape, weight, surface texture, and voice.

This suggests that the medium through which a creative system presents itself is not merely an interface concern. A system’s capacity to be perceived as creative and engaging may depend as much on its material and sensory presentation as on the quality of its generative algorithms. What counts as creative output from an embodied, voice-enabled speculative artifact is fundamentally different from what counts as creative output from a text generation model, even when the underlying language model is identical.

Limitations

Several limitations qualify these findings. The participant sample ($N = 6$) is small and drawn from a single institutional context, limiting generalizability. Each participant interacted with the artifact for a single day; a longitudinal study over weeks or months might reveal different engagement dynamics, including habituation effects or deepened worldbuilding. The system relies on GPT-4o-mini, whose

training data and behavioural tendencies are not fully transparent, making it difficult to determine whether specific creative behaviours are attributable to the prompt architecture or to the model's pre-existing capabilities. Emotional engagement data is self-reported through diary entries and interviews, introducing potential biases related to social desirability.

Finally, the evaluation does not include a comparative control condition, such as a text-only screen chatbot interface. While isolating these modalities could theoretically separate the distinct contributions of form versus text, doing so would fundamentally disrupt the ecological validity required by an experiential futures framework. Because the physical conch explicitly serves to offload cognitive load, it allows users to seamlessly inhabit the speculative context rather than abstractly imagining its parameters through a screen. As such, the study intentionally evaluates the system as an indivisible, multimodal unit.

Future Directions

The most significant design direction emerging from this work concerns the elimination of the screen interface. In its current form, the system relies on a browser-based web application to display transcribed text alongside audio playback. While this was practical for prototyping and evaluation, it reintroduces a screen-mediated layer that partially dilutes the tangible, object-centred experience the physical artifact is intended to produce. Future iterations should move toward a fully embedded implementation in which all computation, audio input, and playback are housed entirely within the conch shell itself, with no external display. This was, in fact, the original design intent of the project, and remains the most coherent realization of the core premise: an artifact that communicates entirely through voice, touch, and physical presence rather than text on a screen.

Concretely, this could be achieved by migrating the system to an embedded single-board computer (such as a Raspberry Pi) running the full pipeline locally or via a minimal wireless API call, with a compact speaker and microphone integrated into the shell cavity. The result would be a self-contained object with no associated screen, no visible interface, and no external device dependency. Such a configuration would more fully support the perception of the conch as an autonomous, quasi-living entity, extending the effect observed in post-interview responses in which participants described the artifact as a relic or being rather than a software system.

Conclusion

This paper has presented Conch(ious), a computationally creative system integrating a persona-constrained LLM with text-to-speech synthesis and a 3D-printed physical artifact, functioning as an autonomous cultural curator from the fictional underwater colony of Amphitopia. Developed through strategic foresight methods, iterative prototyping, and anticipatory ethnographic evaluation, the system produced three significant findings.

First, constraining an LLM within a detailed fictional ontology redirects rather than diminishes its creative capacity.

The Amphitopia persona required analogical reasoning, inversion of perspective, and narrative reframing in every response, producing outputs that participants found generatively rich enough to sustain deepening engagement across repeated interactions.

Second, the consistent emotional trajectory observed across all participants, from curiosity through anxiety to empathy and reflection, demonstrates that provocation is a measurable form of creative output, emerging dynamically from interaction. The critical turn in post-interviews, in which participants shifted from accepting the Conch's narrative to interrogating its ideological assumptions, is particularly notable: the system generated conditions in which users produced their own creative and critical conclusions in response to computationally generated content.

Third, physical embodiment directly shaped how users perceived the system's creative identity. Participants described the conch as a mystical entity whose presence shaped the depth of engagement in ways a screen-based interface would be unlikely to replicate.

Taken together, these findings demonstrate how the boundaries of computational creativity can be extended when generative AI is given not just a task but an identity, a world to inhabit, and a relationship to sustain. Conch(ious) offered a lens through which participants reflected on the fragility of cultural memory and the constructed nature of historical narrative. The system advances a concrete design principle: specificity of constraint, rather than breadth of freedom, can serve as the productive engine of computational creativity, and the medium through which a creative system presents itself materially is integral to its capacity to be perceived as a creative agent.

Acknowledgements

The authors would like to thank the six anonymous participants who generously dedicated their time, spaces, and personal reflections to the anticipatory ethnographic study. We also express our gratitude to New York University Abu Dhabi for providing the lab facilities, equipment, and technical support that made the development of *Conch(ious)* possible.

Author Contributions

Yaakulya Sabbani designed and developed the software architecture, prompt fine tuning, and built the web UI. Ahmad Hafizh led the design of narrative worldbuilding, media production, and integration of primary research into the frameworks. Amina Magomedova contributed to the hardware fabrication, prototyping, and handled the graphic design. The entire study was conducted under the supervision and guidance of Professor Nimrah Syed.

References

- Al-Jayyousi, O. R. 2016. *Islam and Sustainable Development*. Gower.
- Anadol, R. 2023. *Artificial realities: Coral*.
- Boden, M. A. 2004. *The Creative Mind*. Routledge.

- Brierley, G.; Tadaki, M.; Hikuroa, D.; Blue, B.; Šunde, C.; Tunnicliffe, J.; and Salmond, A. 2018. A geomorphic perspective on the rights of the river in aotearoa new zealand. *River Research and Applications* 35(10):1–16.
- Candy, S., and Kornet, K. 2019. Turning foresight inside out: An introduction to ethnographic experiential futures. *Journal of Futures Studies* 23(3):3–22.
- Colton, S.; Pease, A.; and Saunders, R. 2018. Issues of authenticity in autonomously creative systems. 272–279. Proceedings of the Ninth International Conference on Computational Creativity, ICCO 2018.
- Duarte, C. M. 2021. The soundscape of the anthropocene ocean. *Science* 371(6529).
- Dunne, A., and Raby, F. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. MIT press.
- e Cunha, M. P.; Clegg, S. R.; and Kamoche, K. 2012. Improvisation as “real time foresight”. *Futures* 44(3):265–272. Special Issue: Weak Signals.
- ElevenLabs. 2024. Text-to-speech API. Software. Retrieved from <https://elevenlabs.io>.
- Glenn, J. C. 1972. Futurizing teaching vs. futures courses. *Social Science Record*.
- Hines, A., and Bishop, P. 2006. *Thinking about the Future: Guidelines for Strategic Foresight*. Social Technologies.
- Hollan, J.; Hutchins, E.; and Kirsh, D. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction* 7:174–196.
- Holopainen, M., and Toivonen, M. 2012. Weak signals: Ansoff today. *Futures* 44(3):198–205. Special Issue: Weak Signals.
- IPCC. 2023. Climate change 2023: Synthesis report. Technical report, Intergovernmental Panel on Climate Change.
- Ishii, H., and Ullmer, B. 1997. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 234–241. ACM.
- Langton, J. 2019. When the uae was young. *The National News*.
- Li, P.; Yang, J.; Islam, M. A.; and Ren, S. 2025. Making ai less ‘thirsty’. *Commun. ACM* 68(7):54–61.
- Lindley, J.; Sharma, D.; and Potts, R. 2014. Anticipatory ethnography: Design fiction as an input to design ethnography. volume 2014, 237–253.
- Lozano-Hemmer, R. 2013. Voice tunnel.
- Microsoft Research. 2018. Project Natick phase 2. natick.research.microsoft.com.
- Montagu, J. 2018. *The Conch Horn: Shell Trumpets of the World from Prehistory to Today*. Hataf Segol Publications.
- Nimmo, H. A. 1999. The bajau laut: Adaptation, history, and fate in a maritime fishing society of south-eastern sabah. *American Anthropologist* 101:461–462.
- Novaglio, C.; Bax, N.; Boschetti, F.; Emad, G. R.; Frusher, S.; Fullbrook, L.; Hemer, M.; Jennings, S.; van Putten, I.; Robinson, L. M.; Spain, E.; Vince, J.; Voyer, M.; Wood, G.; and Fulton, E. A. 2021. Deep aspirations: towards a sustainable offshore blue economy. *Reviews in Fish Biology and Fisheries* 32.
- Pal, J. S., and Eltahir, E. A. B. 2016. Future temperature in southwest Asia projected to exceed a threshold for human adaptability. *Nature Climate Change* 6(2):197–200.
- Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:76–99.
- Shanahan, M.; McDonell, K.; and Reynolds, L. 2023. Role-play with large language models.
- Subraelu, P. 2021. Sea-level rise and coastal vulnerability: A preliminary assessment of uae coast. *Journal of Coastal Zone Management* 24.
- UNFCCC. 2015. The paris agreement.
- United Nations. 1982. United nations convention on the law of the sea. United Nations Treaty Series, vol. 1833, p. 3.