Gaze to the Stars: AI and Public Art from Personal Affect and Collective Empathy

Behnaz Farahi
Critical Matter Group †
MIT Media Lab
behnaz_f@mit.edu

Sergio Mutis *
Department of Architecture
MIT
smutis@mit.edu

Yaluo Wang *
Graduate School of Design
Harvard University
yaluo_wang@mde.harvard.edu

Suwan Kim *
Department of Architecture
MIT
suwankim@mit.edu

Chenyue "xdd" Dai *
Department of Architecture
MIT
xdd44@mit.edu

Haolei Zhang *
Graduate School of Design
Harvard University
haoleiz@mit.edu

Figure 1: *Gaze to the Stars* projection on the MIT Dome, seen from across the Charles River

Abstract

Gaze to the Stars is an AI-mediated participatory public art installation that transforms the MIT Great Dome into a platform for collective emotional storytelling. 200 Participants engage in reflective conversations with a large language model, their responses encoded as Braille into their iris videos and projected onto the Dome. Using affective narrative embedding, iris segmentation, and emotional clustering, the project visualizes both shared and personal affect. Through *Gaze to the Stars,* we frame AI not as a tool for efficiency, but as a companion—inviting self-reflection and fostering civic empathy within public space.

*These authors contributed equally to this work

†Director of Critical Matter Group: https://www.media.mit.edu/groups/critical-matter/overview/

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Creative AI Track.

1 Rethinking Public Art in the Age of AI

> *In affective economies, emotions do things, and they align individuals with communities—or bodily space with social space—through the very intensity of their attachments.* (Ahmed, 2004)

What if public art could feel, listen, and share? How might it help individuals reconnect with themselves and with one another through lived emotional experiences?

Emotions do not reside solely in the private realm of the self. They extend outward, shaping our actions, informing our decisions, and structuring our relationships. Emotions are the undercurrents of connection and disconnection, as well as inclusion and isolation. They can fuel movements, provoke conflict, or cultivate solidarity. Emotions are shared and relational.

By illustrating *Gaze to the Stars*, we explore how AI-mediated participatory public art can serve as a vehicle for empathy, reflection, and shared emotional experience. Tracing an arc from personal affect to civic connection, we draw on frameworks from affective computing (Picard, 1997), conversational AI, anthropomorphic design (Nass et al., 1994), and empathy-aware storytelling systems. This installation transforms MIT's iconic architectural landmark into an interface for collective storytelling, positioning the participant as the narrator and the AI as a reflective companion.

In this paper, we first examine how *Gaze to the Stars* facilitates self-reflection by leveraging large language models (LLMs) to collect emotionally driven personal narratives. We then explore how these individual affective narratives are transformed into a collective expression of civic empathy, animating the Dome as a living emotional monument.

2 Personal Data Collection Through Conversational AI and Design

As AI becomes more deeply woven into the fabric of everyday life, it is essential to examine how it influences not only what we do, but how we think and feel. Recent studies suggest that large language models like ChatGPT can reduce neural engagement during complex tasks, contributing to a decline in users' cognitive effort and creativity(Georgiou, 2025). This concern echoes Shneiderman's call for Human-Centered AI (HCAI), which promotes systems that amplify human agency (Shneiderman, 2022). *Gaze to the Stars* explores how a semi-agentic AI (one that guides reflection through a structured ReAct workflow) can align with the goals of Human-Centered AI (HCAI). Rather than solving problems, it prompts introspection and supports affective narrative construction.

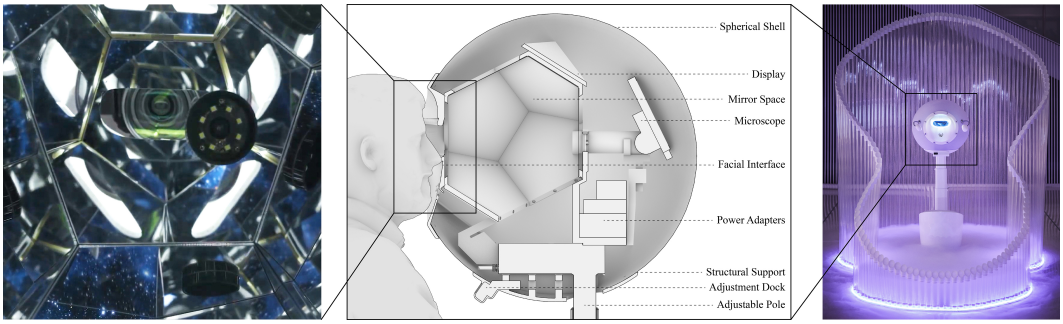

Figure 2: Exterior and interior views of the data collection pod, equipped with a microscope, microphone, and speakers for capturing and responding to participants' stories.

To collect emotional data for *Gaze to the Stars*, we designed an AI-powered pod that gathered over 200 eye videos paired with personal narratives. Participants are invited to step inside the pod, gently rest their forehead against the interface, and gaze into an "infinity room", a mirrored environment that evokes the feeling of stargazing (Figure 2). There, a conversation with the AI begins, while a close-up of their right eye is recorded. The AI is given voice through an anthropomorphized character called the Dome, embodied as a 109-year-old sage. To achieve this, we used ElevenLabs' Multilingual v2 model. Matching the AI's voice to this sage-like character was essential in evoking a meditative, wise

tone. She welcomes each participant and initiates a reflective storytelling process by posing one of four thematic prompts: dreaming, longing, struggle, or emotion [1].

The conversation starts with the Dome introducing herself. To encourage sharing, the Dome begins by recounting anonymized stories she has "heard" over the years, invites participants into a space of mutual empathy. For example, when the theme is *struggle*, participants hear voices saying, "I felt disconnected because I didn't know anybody." Storytelling cultivates empathy by offering listeners access to others' inner lives, bridging social and cultural divides, and establishing trust through shared challenges. Evidence from Minson *et al.* shows that sharing personal narratives containing vulnerability can narrow ideological gaps and substantially increase trust between strangers (Hagmann et al., 2024). These personal narratives were composed by the authors, drawing from MIT student interviews and anonymized public forum posts, then re-scripted and voiced using ElevenLabs.

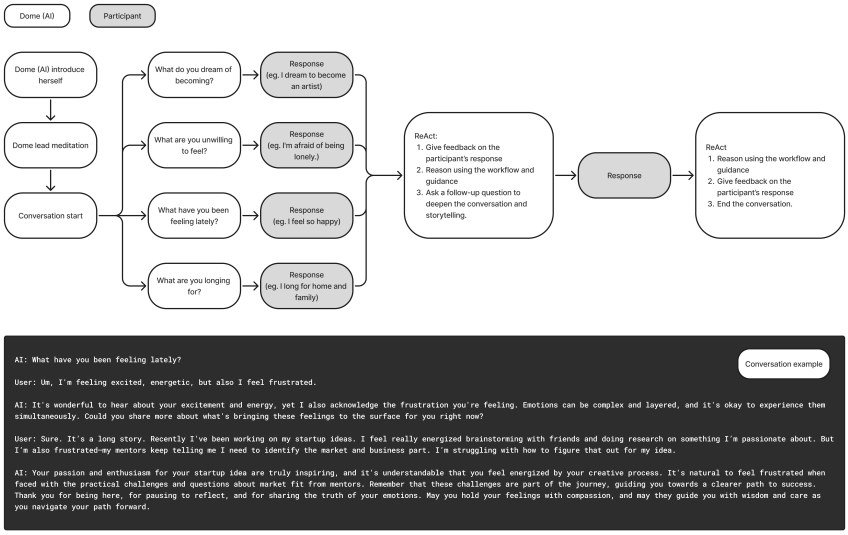

Figure 3: Diagram of the conversational AI for data collection, and an example of user interaction.

The design of participants' interactions (Figure 3) draws inspiration from Marvin Minsky's theory that emotions and cognition are interconnected, operating across six levels of mental activity: instinctive reactions, learned reactions, deliberative thinking, reflective thinking, self-reflective thinking, and self-conscious emotions (Minsky, 2007). The four prompts in *Gaze to the Stars* are carefully crafted to activate distinct layers of cognitive and emotional processing, encouraging participants to reflect on their lived experiences, internal conflicts, and imagined futures. The first prompt, "What is your dream of becoming?", engages self-conscious and self-reflective thinking by inviting participants to articulate long-term aspirations that are tied to personal values and social roles. This evokes hope and pride while prompting reflection on whether these dreams align with one's identity. The second prompt, "What are you longing for?", focuses on learned reactions and reflective appraisal, drawing attention to unmet needs or attachments shaped by past experiences. By reflecting on these longings, participants can distinguish between desires and cravings, contributing to clearer self-understanding. The third prompt, "What are you unwilling to feel?", is inspired by Minsky's distinction between pain and suffering. It surfaces avoidance patterns by targeting learned emotional defenses and encouraging reflective re-framing. This helps participants confront and reinterpret suppressed emotions such as shame or grief, integrating them into a more complete self-narrative. Lastly, the prompt "What have you been feeling lately?" begins with instinctive reactions and moves toward reflective labeling. By focusing on affective states, participants practice metacognition that recognizes and contextualizes emotions so they become part of conscious memory and learning rather than remaining as fleeting experiences. Collectively, these prompts support identity exploration through cognitive-emotional processes, positioning the AI as a guide for introspection rather than a source of answers.

The conversational engine is powered by OpenAI's ChatGPT-4o, accessed via a real-time API. We developed a custom prompt structure to guide each interaction through two rounds of dialogue, with

[1]Each participant was randomly assigned one of four thematic prompts.

each question becoming progressively more reflective. For example, when the theme is struggle, the Dome might begin with: "What are you unwilling to feel?" The AI responds to the participant's tone—whether expressive, hesitant, or silent—with meditative reflections, then invites a deeper story behind their feelings. Each session follows a fixed structure: Two prompts from the Dome and two participant responses were exchanged over a period of three to five minutes.

We collected 200 eye video recordings alongside transcriptions of participant conversations. The conversational AI prompts guide a journey through memory, emotion, and aspiration, positioning the AI as a gentle catalyst for self-reflection. Preliminary observations indicate that participants attributed just enough agency to the AI to feel accompanied, while still retaining ownership of their reflections. We argue that this form of reflective agency presents a compelling design pattern for Human-Centered AI, one that fosters deeper introspection without prescribing answers or promoting cognitive offloading. But what happens when personal reflections become public?

3 Collective Empathy Through Public Art Projection

Gaze to the Stars seeks to extend private narratives into collective resonance. Following Sara Ahmed's theory of affective economies, emotions do not reside solely within individuals but rather accumulate and circulate, "sticking" to objects, places, and bodies. Through this circulation, they generate meaning, shape boundaries, and align individuals with or against communities (Ahmed, 2004). *Gaze to the Stars* engages this terrain by asking: how might AI facilitate not only self-reflection, but shared emotional presence for the community? And what forms of civic empathy might emerge when public architecture becomes a canvas for collective feeling?

Figure 4: (Left) Iris segmentation masks with overlaid pupil positions; (Right) Encoded Braille message embedded within one participant's iris.

To explore these questions, we projected participants' eyes and personal stories onto the MIT Dome, transforming it into a shared emotional landscape. After each participant engaged in an AI-guided conversation, their narrative was summarized and translated into Braille (Figure 4). These messages were then embedded into the iris region of each participant's eye video, turning the eye itself into a vessel of meaning. We designed a precise technical workflow to support this. Using Meta's Segment Anything Model 2 (SAM2) within a custom Jupyter notebook pipeline, we performed iris segmentation on the video frames. Starting from one annotated frame, SAM2 generated a clean binary iris mask, which was propagated across the video using lightweight post-processing (using morphological operations and thresholding). The pupil center was calculated via image moments and normalized, then saved as a JSON file. An adaptive smoothing algorithm reduced jitter while preserving micro-movements, enabling stable, pupil estimation. This approach allowed frame-level alignment overlays synchronized with each eye, even during partial or full eye closure.

These outputs became the substrate for visual encoding. A Unity-based system converted each participant's story summary into Braille, which was rendered in polar coordinates as a point cloud orbiting around the pupil. To generate these messages, we developed a custom program in Unity that converts textual input into Braille, arranged along a polar coordinate grid. Each character is encoded using the standard 6-dot Braille representation and rendered as a point cloud via Unity's Visual Effect Graph. A Graphics Buffer streams spatial data to the GPU in real time, allowing for the fluid animation of thousands of particles forming and orbiting around the central Braille message. A signed distance field (SDF) stamped each Braille dot into a RenderTexture, enabling interactive, gaze-reactive behaviors like collision and modulation, allowing the visual field to respond to the iris's position and blinking.

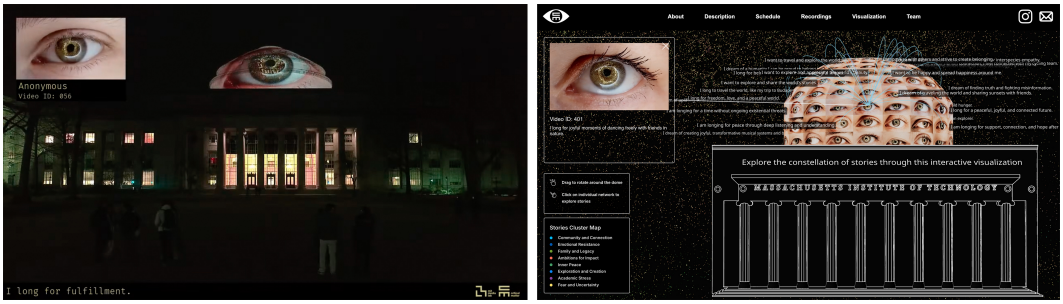

Figure 5: (Left) Livestream of the projection on the MIT Dome, where messages could be decoded; (Right) The interactive data visualization highlights common patterns that emerged across participants.

For the final projected video, we developed a Unity-based workflow pipeline that produces visuals where encoded personal messages and particle fields automatically respond to the viewer's recorded iris movements. For each participant identifier, the system loads a corresponding eye video, iris mask video, textual message, and frame-aligned pupil position data (stored in JSON). During playback, the viewer's gaze position dynamically modulates the visual effect system by updating shader parameters in real time. Our system can render and export these sessions using Unity's Recorder API, generating high-resolution visual outputs for each recorded eye (Critical-Matter-Group, 2025).

Each eye video with encoded messages was composited into a continuous visual sequence. Using Blender and Python, we choreographed the footage with rhythmic camera movement: 10 seconds zoomed in on a single eye, followed by 18 seconds of collective gazes. The final composition was projection-mapped onto the Dome's surface using MadMapper.

The projection transformed the Dome into an affective storytelling vehicle. For onlookers across the city, the giant eyes became portals to individual stories. A custom-built livestream, created using MoviePy and FFmpeg, allowed viewers to decode the emotional messages embedded in each eye (Figure 5, Left). The interface displayed each participant's name or ID, and a narrative summary, enabling audiences both on and off campus to engage with these intimate reflections in real time.

Moreover, the interactive data visualization (Figure 5, Right) invites users to explore each eye in relation to others, uncovering hidden emotional patterns, prompting introspection, and cultivating a sense of shared experience and global empathy. By turning private reflections into a shared emotional landscape, individual emotions become part of a broader constellation of meaning. To do this, we used data clustering techniques and built an interactive digital visualization. Each anonymized narrative was embedded using OpenAI's text-embedding-3-large model and projected into a two-dimensional space using UMAP for interpretability. We then applied HDBSCAN to identify eight thematic clusters, including "community & connection", "family & legacy", and "academic stress". These labeled embeddings were used to calculate pairwise distances, linking each message to at least three nearby neighbors. The results were rendered in a Three.js web interface featuring a virtual MIT Dome, with participant eye videos texture-mapped onto its surface. Hover interactions reveal individual messages, forming an emotional atlas of shared human experience.

Our public art projection and interactive data visualization reveal shared emotional patterns across personal narratives. By uncovering these commonalities, the installation becomes a node of recognition, a space where participants feel seen, affirmed, and emotionally connected to others.

4 Discussions, Limitations & Opportunities

We reflect on key insights gained from the design and deployment of *Gaze to the Stars*.

While this first deployment prioritized emotional authenticity and layered public engagement, future work could extend the system toward real-time feedback, enable a more intuitive message decoding, and establish a more systematic evaluation framework. These additions would offer deeper insight into participants' affective responses and the project's broader impact on civic empathy.

Throughout the project, we noticed the tension of making private emotions public. While participants gave informed consent, the act of externalizing inner life into a shared architectural space

raised important questions. Out of the 200 participants, approximately 88% engaged in recorded conversations, while the rest hesitated—possibly due to ambient noise, unclear cues, or discomfort with trusting a technological interface. These hesitations reveal the delicate nature of emotional vulnerability in semi-public environments. Yet, in an era marked by toxic positivity, where grief, pain, and vulnerability are often overlooked, we believe it is essential to create spaces that acknowledge and honor the full spectrum of human emotion.

We also observed the power of anthropomorphism. Though the pod design did not take on a humanoid form, its AI voice characterization as a 109-year-old sage proved sufficient to encourage trust. Participants frequently attributed human-like understanding to the pod. While it's important to remain critical of systems that simulate human qualities, this subtle anthropomorphic framing seemed to create a safe space for self-reflection. The AI was not treated as an authority or expert, but rather as a reflective companion. Participants could speak freely, share stories, or disengage at any moment, without the emotional labor of reciprocity(Turkle, 2011). This asymmetric dynamic offered a unique opportunity for trust and self-reflection.

Interestingly, the anthropomorphic quality extended beyond the AI voice to the architecture itself. By projecting eyes onto the Dome, the installation transformed a static structure into a living body, watching, sharing, and reflecting. Through this act, the building became a body for communal feeling.

To conclude, rather than imposing collective empathy from the top down, this project allowed it to emerge from the bottom up. Through public projection and interactive data visualization, participants recognized fragments of their own stories in others. In this way, AI and public art converged to foster civic empathy, reminding us that what we feel alone, we often feel together.

After all, we are all stardust.

Acknowledgments and Disclosure of Funding

Gaze to the Stars was developed by the Critical Matter Group at the MIT Media Lab, directed by Behnaz Farahi. The project team included Julian Ceipek, Suwan Kim, Chenyue "XDD" Dai, Sergio Mutis, Frank Cong, Haolei Zhang, Yaluo Wang, Nebus Kitessa, Krystal Jiang, Linda Xue, Yaqi Li, J.D. Hagood, Milin Tunsiricharoengul, Pria Sawhney, and Jiaji Li. Lighting and projection design were realized in collaboration with Agoos D-zines LLC and AVFX. This work was supported by the Massachusetts Institute of Technology, the MIT Media Lab, and Artfinity: An Institute-sponsored event celebrating creativity and community at MIT.

References

Sara Ahmed. 2004. Affective Economies. *Social text* 22, 2 (2004), 117–139.

Critical-Matter-Group. 2025. Gaze to the Stars website. `https://gazetothestars.com`.

Georgios P. Georgiou. 2025. ChatGPT produces more "lazy" thinkers: Evidence of cognitive engagement decline. arXiv:2507.00181 [cs.AI] `https://arxiv.org/abs/2507.00181`

David Hagmann, Julia A. Minson, and Catherine H. Tinsley. 2024. Personal Narratives Build Trust Across Ideological Divides. *Journal of Applied Psychology* 109, 11 (2024), 1693–1715. `https://doi.org/10.1037/apl0001185`

Marvin Minsky. 2007. *The Emotion Machine : Commensense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon Schuster, New York.

Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, Massachusetts, USA) *(CHI '94)*. Association for Computing Machinery, New York, NY, USA, 72–78. `https://doi.org/10.1145/191666.191703`

Rosalind W. Picard. 1997. *Affective computing*. MIT Press, Cambridge, Mass.

Ben Shneiderman. 2022. *Human-centered AI*. Oxford University Press, Oxford, England.

Sherry Turkle. 2011. *Alone Together: Why We Expect More from Technology and Less from Each Other*. Basic Books, New York.