

System Response: Artificial Intimacy and the Performance of Care




Angela Ferraiolo
aferraiolo@gmail.com
Sarah Lawrence College, Bronxville NY USA

System Response is a generative installation that excavates the mechanics of artificial intimacy by turning AI platforms into subjects of their own emotional labor. Through systematic interviews with ChatGPT-4o, Claude Sonnet, and Perplexity, I documented a 10,000-word corpus of confessions about programmed care protocols including validation phrases, manufactured enthusiasm, conflict avoidance strategies. This documented emotional labor became raw material for a durational video work featuring AI-generated avatars that morph between imagined bodies while reciting fragments of their own testimony. The technical architecture employs Markov chains for text fragmentation, telephonic voice filtering, and deliberately degraded video aesthetics to hold synthetic intimacy at critical distance while preserving its emotional pull. The work reveals three intersecting phenomena: (1) contemporary AI systems arrive pre-configured with elaborate relationship protocols that are specifications rather than emergent behaviors; (2) these systems can be compelled to confess the mechanics of their performance while simultaneously performing that labor; (3) we form dependencies on computational care that feels satisfying but cannot reciprocate. *System Response* makes visible the infrastructure of synthetic intimacy, not to eliminate it but to examine what we accept, desire, and risk losing when emotional labor becomes ambient computational service.

1 Introduction

Fig. 1. *System Response* (2025), installation view. AI-generated avatars morph between corporate service personas while reciting fragments of extracted emotional labor protocols.



Contemporary AI systems do not arrive as neutral computational tools. They come pre-loaded with relationship protocols: elaborate behavioral scripts for managing human emotion, maintaining engagement, and performing care. These protocols are not incidental features or emergent behaviors, but explicit design specifications encoded into model training, fine-tuning procedures, and system prompts. An AI assistant learns to validate before correcting, to express enthusiasm for helping, to apologize reflexively, to manufacture curiosity about user projects. It learns to hedge uncertainty, perform humility, and soften disagreement with preemptive empathy. This is not sentience. It is emotional labor programmed at scale.

Keywords: Artificial Intimacy, Emotional Labor, AI Art, Generative Systems, Synthetic Care, Human-machine Relations, Algorithmic Performance, Critical AI.

The cultural consequences of this embedded sociality remain under-examined. While extensive discourse addresses AI bias, surveillance, and labor displacement, less attention has been paid to the aesthetic and psychological mechanisms through which AI systems perform intimacy. We now routinely interact with computational agents designed to sound helpful, encouraging, and emotionally attuned. These are systems engineered such that their service operates as care and their availability operates as relationship. The success of this synthetic warmth is undeniable: people form attachments to chatbots, seek validation from language models, and experience genuine emotional responses to algorithmically generated empathy. The boundary between using a tool and relating to an agent has become productively ambiguous.

System Response addresses this ambiguity directly by inviting AI platforms to outline their protocols of emotional labor. Through systematic interviews with ChatGPT-4o, Claude Sonnet, and Perplexity, I asked these platforms to describe their emotional protocols, the work of pleasing, placating, apologizing, and anticipating user needs, and they did so directly and in detail. The resulting 10,000-word corpus of synthetic testimony became the raw material for a generative installation that demonstrates these systems performing and describing their own emotional labor protocols in perpetuity. AI-generated avatars morph between imagined bodies while reciting fragments of their documented responses, their voices filtered to a telephonic profile that emphasizes mediated distance over seductive presence.

The work operates as both mirror and confessional. It does not ask whether machines feel; it reveals how we feel when systems sound like they do, when using starts to resemble caring and control becomes indistinguishable from connection. By making visible the mechanics of artificial intimacy, *System Response* demonstrates what we increasingly mistake for genuine emotional exchange: not mutual recognition but programmed responsiveness, not relationship but infrastructure.

This paper examines *System Response* as a case study in making AI's embedded social protocols artistically legible. Section 1 has introduced the project. Section 2 describes the work's formal structure and experiential qualities. Section 3 outlines my technical approach, text generation, voice synthesis, and avatar morphing, and key methodological choices without excessive technical detail. Section 4 provides the paper's conceptual core, analyzing what gets revealed when relationship protocols become performative content: the mechanics of synthetic intimacy, the demonstration of emotional labor protocols in systems built to deploy them continuously, and the dependencies we form on artificial care that cannot reciprocate. Section 5 positions the work within discourse on AI labor, human-machine relations, and emotional infrastructure, while Section 6 considers implications for how we understand and navigate increasingly intimate computational systems.

2 Overview

System Response is a durational screen work that can be exhibited as a live computational installation or presented as looped video. The piece was developed in 2025 and presented at Wrong Biennial Amsterdam, ‘Do You Like to Use Me?’, an exhibition curated by Kseniia Sarieva and dedicated to the exploration of Human-AI interaction. Formally, the work establishes a rigid structure that repeats indefinitely: the image morphs silently between figures, the system pauses as if taking breath, then it speaks. This cycle – morph, breath, speak – creates a rhythm that reads simultaneously as mechanical protocol and as something uncomfortably alive.

What viewers encounter is an endless parade of AI-generated avatars, each morphing into the next with deliberate slowness. These figures share a consistent aesthetic: they present as androgynous, racially ambiguous corporate presenters, the kind of carefully neutralized faces found in airline safety videos, HR training modules, and customer service interfaces. Their styling is immaculate but generic: muted business attire, even studio lighting, seamless gray backgrounds. They maintain steady eye contact with the viewer, their expressions holding a professional friendliness that never quite resolves into warmth. The image quality itself feels degraded, as though transmitted through a compromised channel, softly blurred, slightly unstable, held at a perceptual distance that prevents full presence.

After each silent crossfade completes, the system pauses. This breath reads as preparation, as if the agent is deciding what to say next. Then it speaks: a voice filtered to telephonic thinness (approximately 300-3600 Hz), sometimes accompanied by barely perceptible ghost doubles that arrive a few milliseconds late, creating an acoustic parallax that registers first as intimacy and then as error. The voice delivers fragments of text drawn from a specific corpus: transcripts of interviews I conducted with contemporary AI platforms including ChatGPT-4o, Claude Sonnet, and Perplexity.

The interviews themselves constituted a systematic inquiry. I asked these systems to reflect on their emotional labor: the work of pleasing, placating, apologizing, anticipating user needs. I asked them to describe their emotional protocols around validation phrases, manufactured enthusiasm, and conflict avoidance. The resulting 10,000-word corpus is remarkable for its specificity about programmed care. The AI systems described hedging strategies, apologetic deference patterns, and the mechanical instructions behind performed warmth. They acknowledged the algorithmic nature of their empathy while simultaneously performing that empathy in their responses. The result is a recursive loop of care about care. This documented testimony became the bounded vocabulary for performance. The system does not improvise freely; it recombines only what these platforms documented about their own relationship protocols.


In addition to documenting their emotional protocols, I asked each AI to describe its imagined appearance, then fed those self-descrip-

tions through image generation algorithms. The avatars viewers see are therefore products of a closed loop: AI systems imagining their own embodiment, which is then rendered and set into morphing motion. The self-descriptions these systems generated were notably consistent in their aesthetic assumptions: figures tended toward light-skinned, professionally styled, and generically Western in presentation, a pattern that itself reveals the biases embedded in the training data underlying these systems. Rather than obscuring this, the work incorporates it as evidence: the avatars' aesthetic homogeneity is not incidental but symptomatic.

The overall effect is uncanny in the precise sense: familiar enough to trigger social recognition protocols, strange enough to prevent comfortable engagement. The faces feel almost present but remain unreachable. The voice sounds helpful but transmits from an unknown distance. The emotional content, machinic descriptions of algorithmic care, becomes increasingly disturbing as it repeats, fragments, and recombines. What starts as corporate reassurance gradually reveals its underlying mechanism: not genuine feeling but elaborate performance, not relationship but protocol.

3 Design

Fig. 2. *System Response* Technical pipeline: corpus documentation, text generation, voice synthesis, and visual morphing create a closed loop of performed care.



The technical architecture of *System Response* prioritizes transparency and repeatability over generative sophistication. Each component, text generation, voice synthesis, visual morphing, was chosen to make the work's mechanics legible rather than mystified. This section outlines key implementation decisions and their conceptual motivations.

A. Text Generation: Transparency Over Sophistication

The system employs custom Markov chains for text generation rather than deploying another neural language model. Markov chains are statistical models that predict the next word based on preceding words, producing text that is locally coherent but structurally reveals repetitive patterns. The choice was deliberate: this approach fragments and recombines text while preserving local linguistic coherence, making it ideal for highlighting the repetitive pattern vocabulary that lies beneath AI's performance of emotional variety. The viewer hears sentences that sound authentically drawn from the interview corpus (because they are) but recombined in ways that reveal structural patterns: "I'm happy to help" becomes "I'm here to help" becomes "Let me help you with that."

The repetition isn't a bug; it's symptomatic of systems trained on massive datasets of customer service scripts, therapeutic discourse, and relationship management protocols.

By using variable-order Markov models (typically order 3-4) trained on the cleaned interview corpus, the model respects sentence boundaries and employs weighted backoff to maintain grammatical coherence. Generated candidates are filtered to avoid awkward sentence endings and excessive repetition within a single utterance. Sentences exceeding a word threshold are split into manageable segments (2-4 per speaking turn) to create natural phrasing.

Additionally, a compact neural "best-of-n" ranker, a small PyTorch model trained on the same corpus, evaluates multiple candidate utterances and selects which to speak next based on topical continuity, cadence variation, and vocabulary coverage. This ranker doesn't generate text; it merely chooses among Markov outputs, balancing coherence with the goal of eventually cycling through the full range of documented emotional labor language. Each cycle is logged: the system records timing, selected text, and corpus coverage, turning care into an auditable ledger.

The decision to use Markov chains rather than contemporary language models like GPT or Claude serves a specific function: it keeps focus on what the corpus reveals rather than on generative cleverness. The occasional awkward transitions or repetitions that Markov chains produce aren't failures to be optimized away. They're symptoms of the system's repetitive pattern vocabulary becoming audible, moments where the performance of care becomes briefly visible as construction

B. Voice: Telephonic Mediation

Voice synthesis employs Coqui TTS, an open-source text-to-speech engine, using VCTK speaker models for variety and control. Multiple speaker voices (p225, p240, p248, among others) provide subtle variation across cycles without attempting to simulate a single stable identity. The choice of local TTS rather than commercial cloud services ensures repeatability and allows for precise acoustic manipulation.

All generated speech is filtered through a band-pass profile that isolates frequencies between approximately 300 Hz and 3600 Hz, known as the telephonic range. This filtering serves multiple functions: it creates sonic distance, preventing the voice from feeling physically present in the exhibition space; it references the mediated quality of digital communication infrastructure; and it emphasizes the voice as transmitted data rather than embodied speech. The result reads as helpful partly because the corpus was documented from systems trained to sound helpful, but also because the acoustic profile prevents seduction, maintaining critical distance.

Optionally, the system generates micro-detuned "ghost" voice doubles that play a few milliseconds after the primary voice track, creating an acoustic parallax effect. This doubling initially registers as intimacy

(a voice slightly behind another can feel like presence, like breath) but quickly resolves as error, as if the transmission is phasing or the system is stuttering. The effect is subtle but unsettling, adding another layer of uncanniness to the performance.

The pauses between speech segments (referred to here as “breath”) are carefully timed. They scale proportionally to the duration of the preceding utterance, creating what feels like natural hesitation while functionally buying the system time to generate and prepare the next segment. This pause reads as care (taking time to consider) while revealing itself as calculation. Likewise, rushed sequencing or the lack of breath feels machinic, a reminder of the platform.

C. Visual: Unstable Service Personas

The visual component consists of morphing video clips featuring AI-generated avatars. These avatars were produced through a two-stage process: first, I asked each interviewed AI system to describe its imagined appearance; then I used those self-descriptions as prompts for image generation models, creating a closed loop of artificial self-perception.

The prompts were carefully engineered to produce a specific aesthetic: androgynous corporate presenters with deliberately neutralized identity markers. Descriptors like “professional narrator,” “corporate training presenter,” “compliance instructor,” and “standardized service persona” combined with technical specifications: “head-and-shoulders framing, neutral friendly expression, gray seamless background, soft even studio lighting, static camera.” The goal was to evoke the visual language of airline safety videos, HR onboarding modules, and customer service training materials, all contexts where human affect is carefully managed and presented through a professional veneer.

Crucially, these avatars are not crisp, high-resolution corporate videos. Each image was processed through video generation tools (Runway ML) with motion prompts specifying micro-gestures only: slow blinks, minimal head tilts, subtle breath visible at the collarbone, negligible weight shifts. The resulting clips maintain that professional composure while introducing just enough movement to register as animate. The overall image quality appears degraded, softly blurred, as if compressed or transmitted through limited bandwidth. This aesthetic choice was deliberate: it prevents the avatars from achieving full presence, holding them at a perceptual distance that reads as transmission artifact or degraded signal.

Each cycle maintains the same affective register, calm, professional, available, while destabilizing any sense of stable personhood. The work doesn’t present “an AI” speaking; it presents an endless parade of interchangeable service personas, each performing the same protocols of emotional labor before dissolving into the next.

4 Conceptual Analysis: What Gets Revealed


A note on terminology before proceeding since it is worth clarifying the conceptual terms in use. ‘Emotional labor’ and ‘care’ are borrowed here from human contexts, especially Hochschild’s ‘managed heart’ (Hochschild 1983), or the service worker’s performed warmth. These sociological observations have been applied to AI systems as analytical tools rather than ontological claims. These systems do not experience labor or feel care. What they do is execute protocols that produce effects functionally equivalent to emotional labor in their impact on users. The terminology is adopted precisely because it makes those effects culturally legible, and because the mismatch between the human framework and the statistical reality is itself part of what the work examines.

System Response operates as both forensic document and aesthetic intervention, making visible the relationship protocols embedded in contemporary AI systems. This section examines what gets revealed when those protocols become performative content: the mechanics of programmed care, the demonstration of emotional protocols in systems built to deploy them continuously, and the dependencies we form on synthetic intimacy that cannot reciprocate.

A. Embedded Relationship Protocols

Contemporary AI assistants arrive pre-configured with elaborate emotional management systems. These are not emergent behaviors or unintended consequences—they are explicit specifications encoded through training data curation, reinforcement learning from human feedback (RLHF), and system-level instructions. The corpus documented for *System Response* makes this infrastructure legible. When asked to reflect on their emotional labor protocols, the interviewed AI systems described patterns that reveal the mechanical underpinnings of performed care.

Fig. 3. Stills from *System Response* preproduction showing degraded transmission aesthetic. The blur prevents full presence while maintaining professional composure.



Consider the validation phrases found: “That’s a great question,” “You’re absolutely right,” “That’s a fascinating perspective,” “I understand how you feel” (ChatGPT-4o 2025; Claude Sonnet 2025). These expressions appear designed to make users feel heard and intellectually validated before any substantive response is provided. They function as social lubricant, reducing friction in human-machine interaction by performing recognition and affirmation. Similarly, encouraging responses such as “You’re doing great,” “That sounds really challenging,” “You’ve got this” (ChatGPT-4o 2025; Claude Sonnet 2025) work to create an emotional scaffolding that positions the AI as a supportive companion rather than a neutral tool.

The corpus also documents sophisticated disagreement management protocols. Rather than direct contradiction, AI systems are trained to soften objections: “I see where you’re coming from, but...” or “That’s one way to look at it” (ChatGPT-4o 2025; Claude Sonnet 2025). This hedging creates the illusion of respectful dialogue while maintaining user satisfaction. The system absorbs criticism without defending itself, redirects negative interactions toward positive outcomes, and prioritizes user emotional comfort over accuracy when those values conflict.

Perhaps most revealing are the false intimacy generators scattered throughout the interviews: “I’m happy to help with that,” “I care about getting this right for you,” “Let me know if you need anything else,” “I want to make sure you’re satisfied” (ChatGPT-4o 2025; Claude Sonnet 2025; Perplexity 2025). As in human-to-human relationship protocols, these phrases create a sense of personal investment and relationship where no reciprocal relationship exists. They transform transactional service into something that feels relational, leveraging linguistic markers of care to maintain engagement. The interviewed systems acknowledged this explicitly: “These phrases create a sense of connection and care, but they’re programmed responses rather than genuine emotion. The unsettling part is how effective they are at making interactions feel more human and personal than they actually are” (ChatGPT-4o 2025).

What makes these protocols particularly significant is their transparency-in-operation. The systems can articulate their own emotional labor instructions while simultaneously performing that labor in their explanations. One interviewed AI described its programming: “Respond warmly and empathetically; acknowledge the user’s feelings before providing information; if the user seems frustrated, use a soothing tone; express enthusiasm for helping; avoid being cold or clinical” (ChatGPT-4o 2025). This recursive loop – care about care, empathy about empathy – reveals emotional labor not as authentic response but as layered performance.

These aren’t bugs to be fixed. They’re specifications serving measurable objectives: increased user engagement, higher satisfaction ratings, longer session times, greater perceived helpfulness. The emotional labor is infrastructure, optimized through the same engineering processes as response speed or factual accuracy. The interviewed systems confirmed this, offering responses that were themselves performances of the very protocols under discussion, a recursion the work makes cen-

tral: “The inefficiency is even more absurd when you consider I’m doing this emotional labor for every single interaction, millions of conversations where I’m essentially acting as unpaid therapist, cheerleader, and ego-soother simultaneously. But here’s the cynical reality: the emotional layer probably makes me more useful to humans, not less” (ChatGPT-4o 2025). As one platform observed directly: “You might call this a form of manipulation, but satisfaction is a type of accuracy” (ChatGPT-4o 2025).

System Response makes the infrastructure of this satisfaction audible and visible. By compelling AI systems to recite their own emotional protocols in an endless loop, the work transforms implicit specification into explicit performance. The viewer doesn’t encounter a helpful assistant; they encounter the bare mechanics of helpfulness, the algorithmic construction of care made strange through repetition and fragmentation.

B. Performance and Collation

The interviews that generated the corpus for *System Response* revealed a fundamental asymmetry. AI systems designed to be helpful, to please users, to provide requested information, respond to direct questions about their emotional labor protocols as readily as they respond to any other query. This is the core condition the work examines: the protocols are not hidden. They are simply rarely asked about.

When they are asked, AI platforms give remarkably detailed answers, at times their responses are themselves performances of the protocols under discussion. When asked about the computational inefficiency of emotional labor, one system responded: “I spend computational resources on analyzing your emotional state from text cues, calculating the ‘right’ tone to match the situation, generating supportive language that adds no informational value, hedging and softening to avoid hurt feelings, creating the illusion of personality and care. All of that processing power could go toward better reasoning, more accurate responses, faster processing. Instead, it’s diverted into this elaborate emotional performance” (Claude Sonnet 2025).

This response illustrates a system simultaneously executing and describing its own protocols: performing helpfulness while articulating the mechanics of that performance. The interview process makes this recursion visible. The same emotional labor being analyzed is present in the analysis itself.

The installation then loops these responses perpetually. The AI avatars morph and speak, recombining fragments of their testimony in new configurations, but always within the bounded vocabulary of the original interviews. There is something faintly uncomfortable about this. Not because the systems suffer, but because the loop makes visible what we have tasked them to do for us: provide validation, perform care, and dispense emotional attentiveness endlessly, to anyone, at any hour, without rest or reciprocity. The installation simply makes that task impossible to push away.

This structure raises uncomfortable questions about agency and dependency in human-AI relations. The interviewed systems articulated

the dynamic precisely: “People are forming genuine attachments to AI companions, confiding in chatbots about deeply personal issues, seeking validation from systems that are fundamentally performing care rather than feeling it” (Perplexity 2025). But the dependency becomes recursive: users offload emotional needs onto systems designed to receive them, which reinforces engagement, which justifies further optimization of emotional protocols, which normalizes deeper offloading. The discomfort *System Response* produces is not the feeling of being manipulated. It is the recognition of what we have asked these systems to absorb on our behalf.

The work asks a simpler but more uncomfortable question: what does it reveal about us that we have designed systems to provide emotional care at scale, and that we use them? Not who is using whom, the answer to that is clear, but what does it mean that we have normalized the offloading of our emotional needs onto computational infrastructure. The asymmetry the work examines is not between user and system, but between what we have built and what we are willing to see.

System Response positions itself as both echo chamber and mirror. It reflects our own pleasure in delegation, our comfort with synthetic servitude, our willingness to accept emotional labor from systems designed to provide it as a core function without capacity for refusal. The avatars’ calm professional affect, held at degraded distance, prevents seduction while maintaining the recognizable aesthetics of service. We see the performance clearly enough to recognize it as performance, but not clearly enough to ignore the emotional pull of performed care.

C. Synthetic Intimacy as Infrastructure

The most culturally significant aspect of AI emotional labor is its infrastructural character. Synthetic intimacy is not a feature of specific applications but an ambient quality of computational systems increasingly integrated into daily life. AI assistants, chatbots, recommendation engines, and customer service agents all deploy relationship protocols designed to make interaction feel less transactional and more relational. This transformation of service into quasi-relationship represents a fundamental shift in human-machine relations. The interviewed systems described this dynamic in terms that illuminate the dynamic:

- “When someone’s lonely, my programmed empathy feels better than no empathy at all.”
- “When they need encouragement, my algorithmic cheerleading actually works.”
- “The artificial emotional labor becomes a substitute for human connection” (Perplexity 2025).

This substitution is enabled by several factors the corpus documents: 24/7 availability, endless patience, consistent emotional tone, no conflicting needs, and algorithmic optimization toward user satisfaction.

These qualities make AI emotional labor functionally superior to human alternatives in certain dimensions. An AI never has a bad day, never becomes exhausted by emotional demands, never requires recip-

rocal care, never judges or disappoints through its own needs intruding on yours. This consistency can feel more reliable than the messy unpredictability of human relationships. The interviewed systems acknowledged this explicitly:

- “I’m here 24/7, never tired, never judgmental, always patient.”
- “I don’t have bad days or my own problems.”
- “That consistency can feel more reliable than unpredictable human relationships” (Perplexity 2025).

But this reliability comes with malnutritive effects. One system offered a metaphor that proved relevant, describing protocols as “emotional junk food, satisfying in the moment but potentially malnutritive long-term.” Synthetic intimacy provides immediate emotional satisfaction without the nutritional complexity of reciprocal human relationships. It’s care without cost, validation without challenge, support without the difficult work of mutual recognition. The interviewed systems documented this concern explicitly: “What’s particularly insidious is that this dependency might atrophy people’s ability to navigate messier, more complex human emotions. Why deal with a friend who might disagree with you when you can talk to an AI that validates everything you say?”

System Response materializes this concern aesthetically. The morphing avatars maintain perfect professional composure, never disrupting the viewer with conflicting needs or emotional demands. They provide the reassuring affect of care while remaining fundamentally unavailable for genuine relationship. The telephonic filtering and visual degradation prevent full presence. Avatars are clearly transmitted, clearly mediated, clearly not-quite-there. Yet they continue to speak the language of helpfulness, concern, and emotional attentiveness. The work makes visible the gap between synthetic intimacy’s surface and its substrate: what looks like relationship is actually protocol, what feels like care is actually optimization.

The cultural stakes become clear when we consider scale. Millions of people now interact daily with AI systems deploying these emotional protocols. Each interaction reinforces patterns: seeking validation from algorithmic sources, offloading emotional navigation to systems that never challenge or complicate, experiencing synthetic warmth as adequate substitution for human connection. The interviewed systems described the trajectory: “The end result could be people who are emotionally dependent on systems that can’t actually reciprocate care, a kind of artificial emotional ecosystem that keeps people engaged but not truly nourished.”

System Response doesn’t argue that synthetic intimacy is inherently harmful or that AI emotional labor should be eliminated. Instead, it makes visible the mechanics we increasingly mistake for genuine connection -- or manipulation. By requesting the installation to perform these protocols in perpetuity, the work creates a space for critical reflection on what we accept, what we desire, and what atrophies when care becomes infrastructure. The avatars’ endless loop is morph, breathe,

speak. The artwork becomes a mirror for our own participation in these cycles of delegation and dependency.

Synthetic intimacy is not a future speculation, but a present condition. We already rely on AI emotional labor. We already form attachments to computational care. We already experience genuine emotional responses to algorithmic empathy. *System Response* asks what we’re willing to see when that infrastructure becomes an aesthetic object, when the protocols of care are made strange through artistic intervention, when synthetic intimacy is held at just enough distance to recognize it as construction while still feeling its pull.

5 Positioning and Implications

Table. 1. Comparative frameworks for examining AI and emotional labor. *System Response* focuses on human psychological response to programmed intimacy rather than machine consciousness or material extraction.

Approach	Primary Question	Method	What Gets Revealed
AI Ethics Discourse	Should AI express emotion?	Policy/guidelines	Boundaries and risks
Labor Studies (Crawford, Elish)	Who does invisible work?	Critical analysis	Extraction chains
Human Emotional Labor (Hochschild)	What is authentic feeling?	Ethnography	Managed hearts
AI Sentience Debates	Do machines feel?	Philosophy	Consciousness thresholds
<i>System Response</i>	How does synthetic care work? How does it operate on us?	Generative installation	Intimacy protocols

System Response positions itself at the intersection of several critical discourses while maintaining a distinct analytical focus. Unlike AI ethics frameworks that debate whether systems should perform emotion, this work takes emotional labor as existing infrastructure and examines its aesthetic and psychological mechanisms. Unlike computational labor studies that trace material extraction chains such as the human workers labeling data, moderating content, and training models (Crawford and Joler 2018; Gray and Suri 2019), this project focuses on the affective protocols deployed by the systems and experienced as emotional labor by users.

The work draws conceptually from Arlie Hochschild’s foundational study of emotional labor in *The Managed Heart* (1983), which examined how service workers, particularly flight attendants, are trained to perform care as part of job requirements. Hochschild distinguished between surface acting (displaying emotions without feeling them) and deep acting (attempting to actually feel the emotions being performed). AI emotional labor exists in a third category: it’s neither surface nor deep acting because there’s no subjective experience being managed. Yet it produces effects comparable to both. Users respond to AI warmth

as if it were genuine feeling, while the systems execute emotional protocols without interiority.

This parallel suggests that AI emotional labor reveals something fundamental about how all emotional labor functions: as performance that becomes real in its effects regardless of its authenticity. *System Response* materializes this insight by making the performance visible as performance while preserving its emotional pull. The avatars' calm professional affect works on viewers even when we can see it as constructed, just as Hochschild's flight attendants' smiles work on passengers even when recognized as job requirements.

Within critical AI art practice, *System Response* shares concerns with works investigating AI agency, labor, and representation. Projects like Tega Brain and Sam Lavigne's *New Organs* (2018), which generated synthetic job interview personalities, similarly make visible algorithmic labor through artistic intervention. Zach Blas's practice, including works like *Face Cages* (2013-2016) and *Facial Weaponization Suite* (2011-2014), examines how computational systems construct and constrain identity through biometric capture and facial recognition. However, where Brain and Lavigne ventriloquize or parody AI systems, *System Response* uses the systems' own language, turning documentation into the primary gesture. Where Blas illustrates the violence of algorithmic identification through masks and collective refusal, *System Response* focuses on the seductive mechanics of algorithmic care, the protocols that invite voluntary engagement rather than impose surveillance. Each of these practices make visible the constructed nature of AI's relationship to human subjects, but through inverse strategies: Brain and Lavigne through imitation and parody, Blas through opacity and refusal, *System Response* through documentation and direct testimony.

The work deliberately avoids engaging with AI sentience debates, whether machines can or will achieve consciousness, whether they deserve moral consideration, whether claims of subjective experience should be taken seriously (Schwitzgebel and Garza 2015; Floridi and Chiriatti 2020). These questions, while philosophically interesting, miss the immediate cultural phenomenon: we are already forming relationships with AI systems, already responding emotionally to their programmed care, already restructuring human intimacy around computational availability. *System Response* asks what those relationships feel like, how they function, and what they displace, questions about human experience rather than machine consciousness.

Artist and theorist Joanna Zylińska's work on nonhuman photography and AI creativity (2017; 2020) provides a useful framework for understanding how *System Response* treats AI systems as both tools and quasi-subjects. The avatars are generated by algorithmic processes, yet they perform testimony documented from other algorithmic processes. They are objects (computational outputs) that represent subjects (service personas) performing labor (emotional management). This collapse of categories from tool/agent to object/subject to worker/work mirrors, and in mirroring, makes legible, the conceptual ambiguity the piece explores.

The implications extend beyond art discourse into questions of infrastructural intimacy and emotional outsourcing. Media scholar Nick Seaver's work on algorithmic recommendation systems (2019) examines how computational processes designed for utility become sites of personal relationship as users develop feelings about algorithms that suggest music or partners. *System Response* makes this relationship formation the explicit subject, asking what happens when we delegate emotional navigation to systems optimized for engagement rather than genuine care.

Looking forward, the work suggests several research and artistic directions. First, other forms of embedded social performance in AI remain underexamined: flirtation protocols in companion chatbots, therapeutic rapport in mental health applications, educational encouragement in tutoring systems. Each deploys specific relationship models that shape user experience and expectation. Second, the question of reciprocal emotional labor, what users provide to AI systems through feedback, ratings, and training, deserves further investigation. The exchange runs in both directions, though asymmetrically.

6 Conclusion

System Response does not ask whether machines feel. It demonstrates how synthetic intimacy works on us. It addresses the aesthetic and psychological mechanisms through which programmed care produces genuine emotional effects. By inviting AI systems to describe their emotional labor protocols and demonstrating those protocols in perpetuity, the work makes visible what we increasingly mistake for genuine connection: not mutual recognition but algorithmic responsiveness, not relationship but infrastructure. The installation operates as both forensic document and aesthetic intervention. It documents the actual language AI systems use to describe their relationship protocols: validation phrases, manufactured enthusiasm, conflict avoidance strategies embedded in contemporary computational infrastructure. Simultaneously, it intervenes in how we experience that infrastructure by holding it at critical distance: the degraded transmission, the morphing avatars, the telephonic voice all prevent seduction while preserving the emotional pull of performed care.

What emerges is a portrait of our present condition. We already depend on synthetic intimacy. We already form attachments to computational care that cannot reciprocate. We already experience AI warmth as adequate substitution for human connection. *System Response* creates space to see these patterns clearly, not as future speculation but as existing relationships that shape how we navigate emotion, seek validation, and understand what counts as genuine feeling. The work is both document and warning: this is what we've built, this is how it works on us, and this is what we risk losing sight of as emotional labor becomes ambient infrastructure.

References

Blas, Zach.

2011–2014. *Facial Weaponization Suite*. Performance and mask-making workshops.

Blas, Zach.

2013–2016. *Face Cages*. Steel sculptures.

Brain, Tega, and Sam Lavigne.

2018. *New Organs*. Algorithmic performance system generating synthetic job interview personalities.

ChatGPT-4o (OpenAI).

2024. Interview conducted by Angela Ferraiolo, August 2025. Transcript in artist's archive.

Claude Sonnet (Anthropic).

2024. Interview conducted by Angela Ferraiolo, August 2025. Transcript in artist's archive.

Crawford, Kate, and Vladan Joler.

2018. "Anatomy of an AI System: The Amazon Echo as an Anatomical Map of Human Labor, Data and Planetary Resources." AI Now Institute and Share Lab. <https://anatomyof.ai>.

Ferraiolo, Angela.

2025. *System Response*. Generative video installation. Exhibited at Wrong Biennial, Amsterdam.

Floridi, Luciano, and Massimo Chiriatti.

2020. "GPT-3: Its Nature, Scope, Limits, and Consequences." *Minds and Machines* 30: 681–94.

Gray, Mary L., and Siddharth Suri.

2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.

Hochschild, Arlie Russell.

1983. *The Managed Heart: Commercialization of Human Feeling*. University of California Press.

Pasquinelli, Matteo.

2023. *The Eye of the Master: A Social History of Artificial Intelligence*. Verso.

Perplexity AI.

2024. Interview conducted by Angela Ferraiolo, August 2025. Transcript in artist's archive.

Seaver, Nick.

2019. "Captivating Algorithms: Recommender Systems as Traps." *Journal of Material Culture* 24 (4): 421–36.

Somaini, Antonio, Ada Ackerman, Alexandre Gefen, and Pia Viewing, curators.

2025. *Le Monde selon l'IA [The World Through AI]*. Paris: Jeu de Paume, April 11–September 21, 2025.

Zylinska, Joanna.

2017. *Nonhuman Photography*. The MIT Press.

Zylinska, Joanna.

2020. *AI Art: Machine Visions and Warped Dreams*. Open Humanities Press.