The Stochastic Parrot: a physical AI Cohabitant

Ethan Chang
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
echang25@mit.edu

Quincy Kuang
MIT Media Lab
Massachusetts Institute of Technology
Cambridge, MA 02139
quincyku@media.mit.edu

Marcelo Coelho
Department of Architecture, Design Intelligence Lab
Massachusetts Institute of Technology
Cambridge, MA 02139
marceloc@mit.edu

Abstract

In this paper, we discuss a new physical interface towards agentic models. We define AI Cohabitants as new interfaces for agentic models that act with an inherent character, and have the ability to act autonomously without user inputs. An AI cohabitant is more like a roommate or smart house pet, with its own personality and narrative, reversing the subservient dynamic of current AI assistants. Unlike traditional voice agents that require users to actively open an app or issue a command, AI Cohabitants are inherently physical and exist alongside users continuously, occupying space in everyday life with a degree of autonomy. This persistent, ambient presence allows the AI to observe, learn, and subtly participate in the user's world—not by maximizing engagement time or demanding attention, but by developing its own narrative and rhythm of interaction. To explore this concept, we developed a physical embodiment of such a cohabitant: a robotic parrot named the Stochastic Parrot [7]. Through its physicality and contextual framing as a living presence in the user's environment, we observed new patterns of interaction and more nuanced relational dynamics between the AI and its human counterpart. The framing of the AI as a characterful cohabitant, rather than a utilitarian assistant, invites a more spontaneous, expressive, and emotionally textured form of engagement.

1 Introduction

Sci-fi movies since the 1950s have hinted at new forms of human intelligence embedded in computing hardware. From Blade Runner, The Terminator, to Her, these films direct our imagination towards human-form entities that can simulate human-level reasoning and emotional capacities. On the other hand, while we have discourse in creating these new simulated humans, society seems to be more accepting of intelligence in forms that have lower agency and action spaces. From Tomogachi toys in the 1970s, to Furbis and the upcoming Barbie dolls with embedded AI, these affordances to intelligent systems create new interactions that don't directly mimic human interactions.

We define AI Cohabitants as new interfaces for agentic models that act with an inherent character, and have the ability to act autonomously without user inputs. An AI cohabitant is more like your roommate or a smart house pet, having its own personality that revolves around its own identity and narrative. This is a group of new interfaces that doesn't roleplay as an AI assistant, reversing the

Preprint. Under review.

Figure 1: The Stochastic Parrot prototype perched on a user's shoulder.

subservient dynamic of current AI interactions. AI cohabitants are inherently physical. The physical appearances of these agents, being always lit on screen or on object, creates an equal relationship that creates interesting subsequent behaviors. Instead of opening an app to talk to a voice agent when you need to, AI Cohabitants are around you in your life independently of you. This allows the Cohabitant to learn about you, inject its own narrative, and create a closer yet non-anthropomorphic relationship with this non-human companion without the goal of maximizing engagement time. To explore this concept, we created an instance of the AI cohabitant, a physical parrot robot named the Stochastic Parrot. The parrot comes with two servos, a desk mount and shoulder mount, with a camera, a speaker, and a microphone allowing it to engage with the user. The stochastic parrot is in between ChatGPT and your best friend, it is not someone you should find deep empathetic experiences with, but rather a confidant that you trust to experience life with. It is an instance that augments your knowledge and provides feedback in ways you can sculpt while still having its own character.

Table 1: Classification of AI Types by User Dependence and Character Presence

	Without Character	With Character
User Dependent	AI Assistants	AI Companion
User Independent	AI Agents	**AI Cohabitants**

2 Related works

2.1 Digital assistants

The growing accessibility of large language models (LLMs) has enabled everyday users to simulate AI personas for consultation, companionship, and emotional support. Tools like OpenAI's Custom GPT [6] allow users to construct tailored agents through prompt engineering and multimodal inputs, including text and images. Among widely used platforms, Replika [4] prioritizes emotional connection over task efficiency, offering a socially intelligent companion with memory, personality shaping, and empathetic dialogue. Similarly, Character.AI [1] emphasizes persona consistency and improvisational, character-driven interaction, making conversations feel more like engaging with fictional characters than utility tools. Academic work has also explored the motivational and reflective potential of AI personas. Pataranutaporn et al.'s AI-Generated Virtual Instructors [14] show how simulating admired individuals can increase motivation and emotional engagement in learning. In Future You [13], the same author introduces an interactive digital twin that embodies a future version of the user to support self-reflection and long-term goal setting. While these software-based agents are highly customizable and socially expressive, they remain disembodied, lacking the physical presence and real-time nonverbal cues essential to face-to-face social interaction.

2.2 Portable and wearable assistants

Recent advancements in portable AI systems inform our project's design. These devices unlock new modalities for interacting with AI by embedding intelligence into everyday environments. Commercial examples such as Rabbit R1 [3] and Humane AI Pin [2] reimagine human–AI interaction by shifting away from app-based interfaces to persistent, agent-like companions that deliver context-aware responses across a range of tasks. Similarly, the Rewind pendant [5] introduces a novel relationship between user and device by enabling passive recording and retrieval of everyday conversations, reshaping memory practices and attentional habits. Another direction in portable AI explores wearable form factors that integrate more seamlessly with human perception. The Meta Ray-Ban glasses, for instance, merge visual sensing with digital augmentation, supporting a hands-free, camera-assisted interaction model that moves toward human–machine symbiosis. Academic research has proposed more intimate and embodied approaches to AI. AlterEgo by Kapur et al. [11] enables silent, bi-directional communication with AI by detecting neuromuscular signals from internal speech articulators on the skin's surface. This wearable interface supports natural language interaction without spoken voice or overt motion, providing a discreet and private modality for augmenting cognition. Similarly, SixthSense by Mistry et al. [12] explores gestural input combined with wearable projection to support real-time, situated interactions between the user and digital content. While prior systems demonstrate compelling forms of portability and integration, they often frame AI as functional assistants, offering monotone, one-way interactions.

2.3 Social Robots

Animatronics has long captivated both film and consumer product industries for its ability to physically engage users in ways that static objects and disembodied software cannot. Breazeal's foundational work on social robotics argues that for robots to be accepted in human environments, they must engage in social interaction rather than merely perform tasks. Kismet [8], one of her early projects, explored how facial expressions facilitate human–robot interaction. This work later evolved into Jibo [10], a consumer-oriented social robot designed for long-term, emotionally supportive interaction in everyday domestic settings.More recently, Moxie [9] represents a new generation of socially interactive robots, using affective AI to engage children in themed weekly activities grounded in behavioral science, including principles from cognitive behavioral therapy (CBT) and applied behavior analysis (ABA). Studies have shown that physically embodied agents like Moxie significantly outperform virtual agents in fostering engagement, empathy, and self-regulation skills in children. However, these animatronic-style robots are typically stationary, designed to reside in fixed locations within the home. Their lack of mobility limits their ability to collect contextual data across varied settings and restricts opportunities for continuous, real-time interaction as users move through their environments. We mapped the various existing ai agent product in Figure 2.

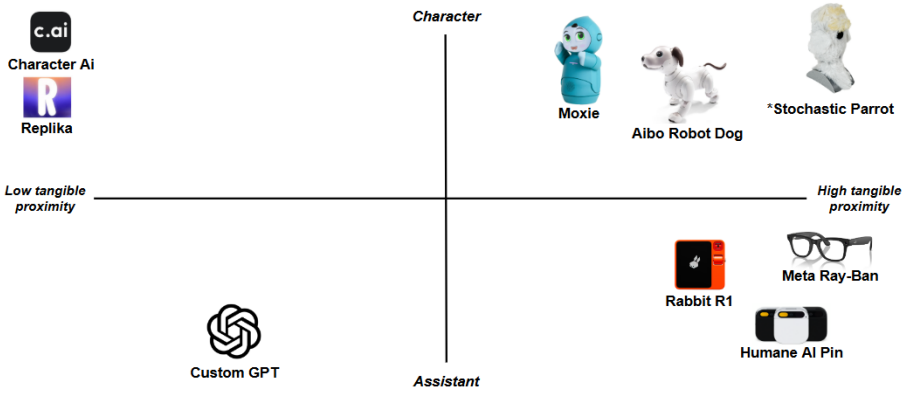

Figure 2: Positioning of various ai agent products on the market. Note that "high tangible proximity" refers to both the physical nature of the ai agent and their accessibility to the user.

3 Design

Users interact with the stochastic parrot in ways similar to how people treat pet parrots. Users can talk, pet, or bring the parrot around and let it rest on the user's shoulders. The stochastic parrot comprises a speaker, a microphone, and a camera as input/output interfaces. The parrot has two degrees of freedom from two servos and a mount that allows users to put it on a shoulder or the desk. The parrot backend is from the ChatGPT Realtime model with a voice filter. The voice model is presented as a cheerful parrot through a system prompt, along with its backstory. The model also allows tool calls to take a picture of the current view, retrieve memory or save memory from a simple memory structure, or move the body of the parrot in certain patterns based on the context of the input voice.

The physical form of the Stochastic Parrot is intentionally non-human: an anamorphic bird with fuzzy fur and an expressive silhouette (see figure 1). We intentionally avoid humanoid representations in favor of animal-based characters, which carry culturally ingrained associations of companionship, trust, and emotional safety. By animalizing the AI agent, we open the possibility for users to relate to it as a distinct character rather than a utilitarian device. Our system explores a different trajectory: an AI agent that is not only wearable and portable, but also socially expressive and character-driven. We hypothesize that this design choice fosters richer, more frequent, and more emotionally resonant interactions.

The primary design contribution of the Stochastic Parrot is the articulation of a novel framing for a wearable, individualistic physical agent interface, defined by two key principles: (1) establishing a persistent and embodied presence through its portable, physical form, and (2) cultivating an autonomous persona that enables the parrot to interact with the world independently of the user's presence.

3.1 Physical presence and portability of an AI cohabitant

We design AI cohabitants to have a strong physical presence by making robots portable and capable of accompanying users throughout their day. We believe that portability is a key component of embodiment. A physically present AI agent can tailor its behavior to the specification of context, which fixed-location systems inherently lack.

Moreover, conversation is often spontaneous and context-dependent. The static nature of most social robots limits interaction to pre-planned or intentional moments, reducing opportunities for serendipitous, on-a-whim exchanges. We hypothesize that portable AI agents can better facilitate these spontaneous interactions, fostering a more natural and fluid conversational flow. In our observations, placing the parrot on the user's shoulder reduced interaction friction. The user could simply turn their head and speak to the parrot, enabling immediate, in-context exchanges without breaking the flow of the moment.

Users can bring the parrot around and show it around the kitchen, classroom, or the outside world when needed. The parrot is also independent in the way that while remaining on the table, it can look around and observe the environment, or observe interactions of the user with other people. We embedded simplified physical motion within the interaction flow, creating a more intuitive rhythm of exchange that mirrors human conversational dynamics. The movement is intentionally restrained; it is sufficient to convey responsiveness without being overly expressive or anthropomorphic. This balance allows the parrot to feel alive and attentive, while maintaining its identity as a mechanical companion.

3.2 Non-human and non-sychopathic nature of AI cohabitants

Conventional AI systems often struggle to elicit meaningful emotional resonance from users. We believe this limitation stems, in part, from the inherently subservient design of most AI platforms, which reinforces an imbalanced relationship between the human and the agent. When an AI behaves purely as a tool—obedient, expressionless, and overly accommodating—it becomes difficult for users to form authentic emotional connections.

Our design leverages a characterful framing allowing the Stochastic Parrot behaves less like a tool and more like a non-human cohabitant. Its refusal to overly compromise or comply is not perceived as malfunction, but rather as personality. Much like how we might describe someone as having a

"punchy" or "fiery" temperament, the parrot's less predictable responses lend it a sense of presence that users can react to, reflect on, and potentially bond with.

We deliberately designed the parrot's persona to be empathetic yet rational, occupying a space between emotional engagement and functional support. In our initial use case, the parrot acts as a motivational coach—a role that straddles the boundary between emotional interaction and utility. Motivational dialogue offers a context where emotionally evocative language can still feel purposeful, providing the AI space to trigger affective responses without overstepping into therapeutic claims. The personality of the parrot is continuously refined through prompt engineering as shown in figure 3, maintaining consistency in tone, temperament, and emotional framing. Additionally, it features a goal-tracking mechanism that stores the user's long-term objective, allowing the parrot to periodically remind the user and provide motivational nudges. This goal system, combined with a persistent memory architecture, enables the parrot to track progress over time and maintain a coherent, evolving character.

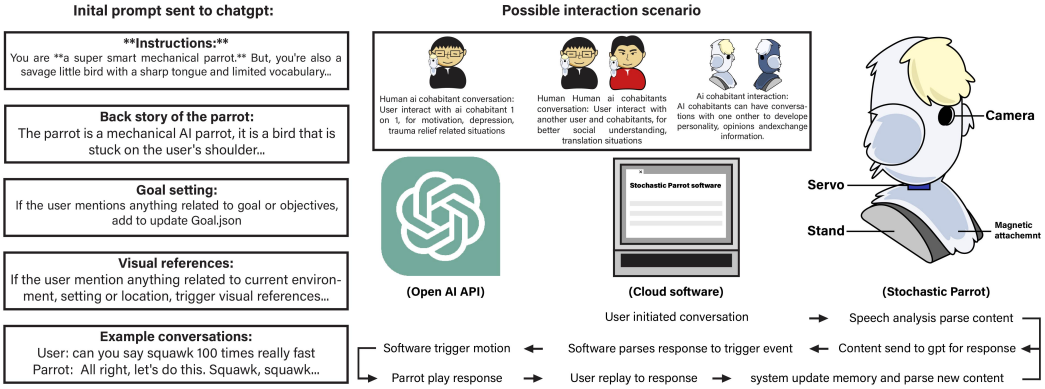

Figure 3: Stochastic parrot prompt engineering and workflow chart

4 Preliminary user testing

We conducted an informal user testing session with seven participants to evaluate the first prototype of the Stochastic Parrot. The primary goal of this preliminary study was to explore moments of unexpected or meaningful interaction that might emerge between users and an AI cohabitant. Through casual observation and open-ended feedback, we aimed to surface early signals of engagement and character perception.

In one notable exchange, a participant asked the parrot a playful question: "What would you want to eat right now?" The parrot replied, "I can't taste anything because I'm a mechanical parrot, but if I could, I'd want to taste success." The participant remarked that she was surprised by the humor in the response and found it unexpectedly charming. We interpreted this moment as a successful outcome of the system's non-sycophantic prompting strategy and the consistent reinforcement of the parrot's motivational persona. Such moments suggest that even early-stage interactions with an "independent" AI agent can evoke delight and emotional resonance.

In another exchange, a participant asked the parrot to repeat "sqwak" 100 times in a very fast speed. The parrot simple refuses the ask and replied by saying "saying sqwak that much isn't really productive, let's focus on something more important in your life instead". We find this refusal from the parrot gives it an interesting character and establishes mutual respect.

The stochastic parrot also provides interesting interactions in a social context. In one occasion, a participant attempted to read a book with the parrot, we observed the dialogue with the parrot became frictionless to provide the parrot the context understanding. The authors also wore the parrot on the campus in public, and one passersby's immediate reaction is to pet the parrot on the author's shoulder, while the parrot chimed in on the conversations. We interpret this interaction as a positive signal towards having the parrot in new social interactions.

5 Discussion and Future Work

AI cohabitants offer an alternative framing for integrating physical AI objects into everyday life. While they may not be as functionally capable as humanoid robots or task-oriented coding agents, they provide a different design path: one that avoids highly attention-seeking or sycophantic robotic forms, as well as overtly human-like embodiments that attempt to replicate the emotional cues of human interaction. Many AI products today aim to simulate human voices, expressions, or presence; the Stochastic Parrot explores how companionship can emerge from new interaction paradigms rather than replicating existing human dynamics. We argue that AI cohabitants can support specific forms of emotional value, though their scope is constrained by the current capabilities of large language models. Some types of emotional support—particularly those that require deep empathy or nuanced emotional understanding—remain beyond the reach of non-sycophantic agents. However, in cases where the goal is to provide a light motivational "push" rather than deep emotional processing, a non-sycophantic AI cohabitant can be both effective and socially acceptable. For example, motivation often requires reinforcement rather than emotional consolation, making it a suitable domain for this design approach. Conversely, in emotionally fragile situations—such as grief or trauma—it would be unwise to rely on a non-sycophantic agent, as the user's emotional state may be too vulnerable to withstand challenge. Based on our early experiences, we identify several scenarios where non-sycophantic AI cohabitants can provide value without overstepping emotional boundaries:

- Lightly pointing out personality flaws in a constructive way.
- Offering encouragement during post-breakup recovery.
- Providing environmental or contextual reminders (e.g., recalling details about a person the user has forgotten).
- Supporting adherence to self-imposed commitments, such as diet goals.
- Motivating the user to get out of bed.
- Helping to shift short-term negative mood swings.

5.1 Directions for Future Work

To advance the design and understanding of AI cohabitants, we see several important research directions:

1. **Autonomous environmental awareness** — Investigating how AI cohabitants can observe and interpret their surroundings without direct user input, enabling more contextually relevant and timely interactions.
2. **Conversation initiation strategies** — Exploring under what conditions the cohabitant should initiate interaction to balance presence with respect for user attention and autonomy.
3. **Interaction preference contexts** — Understanding scenarios in which users prefer engagement with an AI cohabitant over human interaction, and how this preference varies across cultural, situational, and personality factors.

6 Conclusion

The Stochastic Parrot demonstrates how AI cohabitants can occupy a unique space in human–AI interaction, avoiding human mimicry and excessive sycophancy in favor of a consistent, independent persona embedded in a portable, physically present form. Our preliminary observations show that such agents can enable spontaneous, in-context exchanges, provide light emotional reinforcement, and create new forms of companionship beyond traditional human–AI paradigms. While current large language models limit the depth of empathy these systems can offer, they excel in domains where gentle motivation, contextual reminders, and playful interaction enhance daily life. By treating AI not as a mirror of humanity but as an independent character in our shared spaces, we open the door to entirely new ways of living with machines.

References

[1] Character.ai, 2025. Persona-consistent chatbot platform.

[2] Humane ai pin, 2025. Wearable AI device with agent-like interaction.

[3] Rabbit r1, 2025. Context-aware, persistent AI companion device.

[4] Replika, 2025. Online emotional-first AI companion.

[5] Rewind pendant, 2025. Passive voice recording and memory retrieval device.

[6] Sagiv Antebi, Noam Azulay, Edan Habler, Ben Ganon, Asaf Shabtai, and Yuval Elovici. Gpt in sheep's clothing: The risk of customized gpts. *arXiv*, 2024.

[7] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? . In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383097. doi: 10.1145/3442188.3445922. URL https://doi.org/10.1145/3442188.3445922.

[8] Cynthia Breazeal. *Designing Sociable Robots*. MIT Press, 2002.

[9] Nikki Hurst, Caitlyn Clabaugh, Rachel Baynes, Jeff Cohn, Donna Mitroff, and Stefan Scherer. Social and emotional skills training with embodied moxie. *arXiv*, 2020.

[10] S Jeong et al. Deploying a robotic positive psychology coach to improve college .... *arXiv*, 2020.

[11] A. Kapur et al. Alterego: Silent speech interface using neuromuscular signals. In *UbiComp*, 2018.

[12] P. Mistry et al. Sixthsense: A wearable gestural interface. In *SIGGRAPH*, 2009.

[13] Pok Pataranutaporn et al. Future you. *arXiv*, 2024.

[14] Pok Pataranutaporn et al. Ai-generated virtual instructors. *arXiv*, 2024.