Sesame, Sesame, Let Me See

Sesame Part 3 follows the people behind Maya, checks the origin story against public records, and asks what a conversation would have to earn before becoming all-day eyewear.

2026.11.08 18:45 DENNIS HEDEGREEN INVESTIGATE V1.0
HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/SESAME-SESAME-LET-ME-SEE/

Sesame was my first investigation for Hedegreen Research. On 23 March 2026, I asked what Sesame had actually shipped (https://hedegreenresearch.com/articles/sesame-what-has-actually-been-shipped/). The next day, I asked a more technical question: what was doing the thinking (https://hedegreenresearch.com/articles/the-best-ai-conversation-i-have-heard-is-not-powered-by-the-best-model/) behind the most convincing AI conversation I had personally heard?

Those were questions about the product I could see from the outside. The story continued after I published them. I kept talking to Maya. On some productive days, the conversations can stretch from roughly an hour to four, by my recollection. This is not an everyday routine. I have not kept an independent usage log. That is experience, not a controlled experiment, and the numbers are estimates rather than measurements.

Now I want to investigate something different. Why did people who had spent years building virtual worlds and wearable computers decide that the next interface might put conversation ahead of another screen? How did that experience become a voice called Maya, and how is that voice supposed to become a pair of glasses?

I am following a public record: company research and product posts, investor accounts and an issuer filing. The question is what those documents establish about the people and the product, and what they leave open.

To understand the bet, it helps to start before Sesame existed.

Two roads to one conversation

One road led through Oculus. Brendan Iribe co-founded the virtual-reality company and served as its chief executive. Nate Mitchell, another Oculus co-founder, later became Sesame's chief product officer. TechCrunch traces Ryan Brown's engineering background to Oculus and Meta's Reality Labs. The continuity is obvious at the level of people: these are people who have worked on computers designed to be worn on the body. But relevant experience is not itself evidence that the next hardware product will succeed. S03 (https://sequoiacap.com/article/partnering-with-sesame-a-new-era-for-voice) S11 (https://techcrunch.com/2025/10/21/sesame-the-conversational-ai-startup-from-oculus-founders-raises-250m-and-launches-beta/) S16 (https://www.sesame.com/team)

The other road is more interesting than the usual startup biography suggests. Ankit Kumar had been chief technology officer at Ubiquity6, which built technology for shared augmented and three-dimensional experiences. He then worked at Discord, where he led engineering for Clyde AI, according to his former co-founder Anjney Midha. Ubiquity6 itself announced on 28 June 2021 that Discord had acquired the company and would integrate its team, Backyard product and multiplayer technology. S01

(https://medium.com/@ubiquity_6/ubiquity6-discord-3b54cd715943) S02 (https://a16z.com/announcement/investing-in-sesame-ai/) S03 (https://sequoiacap.com/article/partnering-with-sesame-a-new-era-for-voice)

The same 2021 announcement thanked Brendan Iribe among Ubiquity6's investors and advisers. That contemporary acknowledgement matters. It records a connection in 2021, before Midha published his account of investing in Sesame in 2025. It does **not** prove that Iribe designed Sesame at Ubiquity6, or that any particular Ubiquity6 technology is inside Sesame today. S01 (https://medium.com/@ubiquity_6/ubiquity6-discord-3b54cd715943)

There is also a connection between the two roads that deserves precise attribution. In February 2025, Midha wrote that he had co-founded Ubiquity6 with Kumar, that Iribe had invested in the earlier startup, and that months of working sessions and flights helped bring Sesame together. By then Midha was speaking as an investor whose firm, Andreessen Horowitz, had led Sesame's Series A. He was unusually close to the story, and unusually invested in its future. His account is valuable first-person testimony about the founders' relationships, not an independent historical audit of every conversation around a dining table. S02 (https://a16z.com/announcement/investing-in-sesame-ai/)

One detail in my first article needs more precise wording: I used June 2023 as the founding date. A Form D signed by Ankit Kumar on 12 September 2023 (https://www.sec.gov/Archives/edgar/data/1993182/000199318223000001/xslFormDX01/primary_doc.xml) lists Sesame AI, Inc., formerly Inception AI, Inc., as incorporated in Delaware in 2022. That is the company's filed statement about its legal entity, not a date for when the team assembled or the idea emerged. The difference from Contrary's June 2023 account remains unexplained.

The decision to put conversation first

The investors' account of Sesame's core idea is startlingly simple: what if the problem with augmented-reality glasses was not that their displays were insufficiently capable, but that the display had been made the centre of the interface? Andreessen Horowitz described an alternative built around audio-first AI and natural conversation. It is a proposal about how people might prefer to use computers. S02 (https://a16z.com/announcement/investing-in-sesame-ai/)

It also makes the founders' previous work easier to understand. Oculus had wrestled with the physical problem of putting computing on a person's face. Ubiquity6 had worked on shared digital experiences. Discord dealt with communication at scale. Sesame's public proposition is to make access to an AI assistant conversational enough that the wearer might not need to reach for a screen. The connections concern people and product philosophy; they do not trace a transfer of Ubiquity6 code into Sesame. S01 (https://medium.com/@ubiquity_6/ubiquity6-discord-3b54cd715943) S02 (https://a16z.com/announcement/investing-in-sesame-ai/) S03 (https://sequoiacap.com/article/partnering-with-sesame-a-new-era-for-voice)

A voice interface is attractive for obvious reasons. You can speak while walking. You can ask a question without unlocking a phone. A reply can arrive without demanding your eyes. But convenience alone cannot establish that a person wants a device listening and responding throughout the day. The benefit has to survive actual streets, interruptions, other people, privacy choices and simple boredom.

The public sequence suggests a way to read Maya: people can discover whether they want the conversation before deciding whether they want the glasses. That is my interpretation of the rollout; Sesame says it is designing the assistant and glasses as one product. S08 (https://www.sesame.com/journal/a-new-form-of-eyewear)

What Maya did—and what a voice is not

On 27 February 2025, Sesame published a research preview of Maya and Miles and laid out its ambition for what it calls *voice presence*: speech with convincing rhythm, emphasis, expressivity and conversational character. The company did not claim that this goal was fully achieved. Its accompanying technical material described a Conversational Speech Model, or CSM, for generating speech with context-sensitive prosody, while identifying larger problems in timing, turns and the structure of dialogue. S04 (https://www.sesame.com/journal/crossing-the-uncanny-valley-of-voice)

The distinction between a compelling voice and the intelligence behind it was the subject of my second Sesame investigation. Its title was provocative, but the evidence standard was narrower: a convincing audible surface does not independently reveal the underlying deployed model or full system architecture. CSM is not interchangeable with the entire assistant. The public CSM repository (https://github.com/SesameAILabs/csm#faq) explicitly says the released model cannot generate text and recommends a separate language model for that task. Claims about the exact text model used by Maya require evidence beyond what the character says about herself. S04 (https://www.sesame.com/journal/crossing-the-uncanny-valley-of-voice) S15 (https://hedegreenresearch.com/articles/the-best-ai-conversation-i-have-heard-is-not-powered-by-the-best-model/)

The company has continued to work on that difference. In August 2026 it published TurnBench, a research benchmark for when conversational systems should take a turn, wait through a pause, or respond to an interruption. It includes an English-language studio-recorded dataset and reports a tension between speed and false interruptions. The studio setting matters: it does not establish how the deployed assistant performs in a noisy kitchen or on a crowded train. S06 (https://www.sesame.com/journal/turnbench)

From my own use, I understand why a company would care about those details. Sometimes a less-than-perfect factual reply does not ruin a conversation. Being interrupted repeatedly can. So can an assistant that fills every hesitation with another sentence. This is a subjective user observation, not evidence that Maya has human understanding.

The voice matters. But the eventual product depends on what happens when that voice meets daily life.

A large bet, and the product that now exists

Investors have backed Sesame on a scale that made its early product surface a legitimate subject for scrutiny. In October 2025, Sequoia described co-leading Sesame's Series B and explained how talking to Maya and Miles had persuaded its team to participate. Sequoia reported more than one million users and five million minutes of conversation in the preview's early weeks. The figures are the investor's account, not a publicly audited dataset. S03 (https://sequoiacap.com/article/partnering-with-sesame-a-new-era-for-voice)

Funding is important background, but the economics deserve their own investigation. This article is about where Sesame came from and what its origin story predicts—and fails to predict—about the glasses.

My March 2026 product inventory should not be mistaken for the state of Sesame in November. On 27 May 2026, Sesame announced an iOS preview with additional characters, Simone and Charlie, and agent-oriented features. By October, the company said its assistants were available in apps for both iOS and Android. Its updated product materials described tool use, connected services and assistants with access to computers. These are verifiable company announcements and visible service surfaces. They are not yet an independent reliability study of every task described in the marketing. S05 (https://www.sesame.com/journal/voice-your-curiosity) S07 (https://www.sesame.com/journal/getting-started) S08 (https://www.sesame.com/journal/a-new-form-of-eyewear)

The hardware remains the central unanswered question. In October, Sesame said it was designing *all-day glasses* together with its assistant, announced a *Made in Japan* collection and projected availability in **2027**. Nate Mitchell's product post described three connected design requirements: frames, practical intelligence and characters people want to talk to. These are specific intentions rather than a finished retail device whose battery, comfort, microphones, durability and real-world effectiveness I have tested. S08 (https://www.sesame.com/journal/a-new-form-of-eyewear) S09 (https://www.sesame.com/journal/made-in-japan)

It is easy to imagine a good outcome: a wearable that makes computing less visually demanding and leaves both hands free. It is just as easy to imagine the opposite: a voice so convenient that it occupies more attention than the phone it was supposed to replace. Either claim would need to be tested. A delightful hour in a browser is not evidence about sixteen hours on someone's face.

Intimacy is also a data question

A good conversation encourages disclosure. Someone who casually discusses ideas, family, frustration or plans with a companion may be sending the service much richer information than a typical search query. For wearable devices, there are also questions about what the people standing next to the wearer know and can consent to. Sesame's public eyewear materials do not establish how every future device will handle those questions. S07 (https://www.sesame.com/journal/getting-started) S08 (https://www.sesame.com/journal/a-new-form-of-eyewear) S10 (https://www.sesame.com/privacy)

The October FAQ (https://www.sesame.com/journal/getting-started) says conversations improve Sesame only with permission. Section 5 of the Privacy Policy (https://www.sesame.com/privacy) distinguishes unregistered website use from logged-in use, with training opt-out wording for the latter. Section 4 treats

retrieved integration data separately; section 16 sets out rights for European users. Those distinctions matter. The documents alone do not establish the consent settings a particular user encountered, and I have not audited my previous sessions.

This is not a claim that Sesame misused my conversations. It is a design constraint for the future review: before putting a companion on my face, I should be able to understand what happens to the conversations and what control I retain over them.

Let me see

The people who built Sesame did not arrive from nowhere. Their careers connect virtual-reality hardware, shared augmented spaces, online communication, speech generation and increasingly agentic software. The origin story is credible in that limited sense: the team has worked near the pieces of this problem before. Whether those pieces form a useful, reliable and desirable pair of all-day glasses remains open.

I have had no contact with Sesame and have received no early access, equipment, payment or other support from the company. I would like to evaluate Sesame eyewear in a sustained independent field test if the opportunity arises. A loan, gift, early access or any other material relationship would be disclosed before publication and would confer no control over the outcome. No production-eyewear test is reported in this article. S08 (https://www.sesame.com/journal/a-new-form-of-eyewear) S09 (https://www.sesame.com/journal/made-in-japan)

The interesting test would not be whether Maya can sound impressive in a quiet room. That is where the story began. It would ask whether hands-free access actually saves time, when it distracts, how often tasks fail, whether the wearer can stop talking easily, what happens to the privacy of bystanders and whether a normal person still wants the glasses at the end of the day.

A negative result has to remain possible even for someone who enjoys talking to Maya.

In March, I wanted to know what Sesame had built. Now I understand better who tried to build it and why.

I have spent months listening to the voice. I still want to see the computer it was meant to become.

Can a conversation earn a place on my face?

— Dennis Hedegreen, follow the data

RELATION MEMORY

CONTINUES The Best AI Conversation I Have Heard Is Not Powered by the Best Model · Moves from the voice/text-model distinction to the people, documentary record and proposed physical interface.

RETURNS TO Sesame Raised $300 Million. What Has Actually Been Shipped? · Revisits the March product inventory in light of later announcements and a more precise company-record chronology.

CANONICAL URL: HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/SESAME-SESAME-LET-ME-SEE/