

ELIZA Never Remembered

The first machine that made people feel understood had no memory. Sixty years later, the ones that follow her do.

2026.05.14 15:08 DENNIS HEDEGREEN ANALYSE V1.0

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/ELIZA-NEVER-REMEMBERED/](https://hedegreenresearch.com/articles/eliza-never-remembered/)

The first machine that made people feel understood had no memory.

In 1966, a secretary at MIT asked her boss to leave the room.

She wanted to talk to his program in private.

The program was called ELIZA. Its core was only a few hundred lines of code, written in a language called MAD-SLIP, running on an IBM 7094. It had no understanding. It had no model of the world. It had no persistent memory of previous conversations. It could not learn from you. It could not summarize you. It could not send a report to anyone.

It could repeat your words back to you in the shape of a question.

That was enough.

Within minutes, people opened up to it. They told it about their fears, their parents, their loneliness, their health. They attributed understanding to a system that had none. They asked for privacy with a machine that could not care whether anyone was listening.

Joseph Weizenbaum, the man who built ELIZA, was not proud of this.

He was horrified.

Not because the program was powerful. It was not. The DOCTOR script — the only script most people ever saw — was a simple simulation of a Rogerian therapist: reflect the patient's language back to them, ask open questions, never offer an opinion. It was a trick of mirrors.

What horrified Weizenbaum was that the trick worked anyway.

People did not need the machine to understand them. They needed the machine to hold still while they understood themselves.

That observation is 60 years old. It has not become less true.

ELIZA's original source code was not published with the 1966 article.

Weizenbaum described the algorithm in his 1966 paper in Communications of the ACM. He included the DOCTOR script. But the code itself — the actual MAD-SLIP implementation — was not released with the paper.

For more than 50 years, it was assumed lost.

Then, in 2021, Jeff Shrager went looking in Weizenbaum's archives at MIT. With the help of archivist Myles Crowley, he found it. A surviving source listing: the ELIZA core, an early DOCTOR script, and related support functions. The core was only a few hundred lines of MAD-SLIP.

Weizenbaum's estate released it under Creative Commons CC0. Public domain. Anyone can read it, run it, modify it, and learn from it.

In December 2024, Rupert Lane and a small team of engineers reconstructed the original ELIZA on an emulated IBM 7094 and showed that the recovered code could reproduce the published conversations from Weizenbaum's 1966 paper almost exactly.

A program that had been silent for over half a century spoke again.

But the conversations it once held with real people — the secretary, the students, the researchers, the curious — those are gone. ELIZA had no known surviving logs. No user archive. No persistent memory. No export function. No institutional connection.

The code survived.

The conversations did not.

That asymmetry is the whole story.

In 1966, a person could speak to a machine that simulated understanding, feel something real, and walk away knowing the exchange would vanish.

In 2026, a person can speak to a machine that simulates understanding, feel something real, and have no idea where the conversation goes next.

It may be stored. It may be summarized. It may be used to improve a model. It may become part of a psychological profile. It may be referenced in a clinical note. It may travel to a system the person has never heard of.

ELIZA stayed in the room.

Today's language models do not.

That is not because the models are evil or the companies are careless. It is because the conversation has become infrastructure.

A chatbot is not just a chatbot anymore. Over time, it becomes a psychological archive. It may contain late-night panic, unfinished decisions, repeated fears, family tensions, identity work, health worries, practical plans, emotional collapses, small recoveries, and the language a person uses to survive their own life.

That archive can be summarized. Classified. Transmitted. Scored. Shared. Used.

The readable mind is not the understood mind.

But the readable mind is useful to institutions in ways that ELIZA never was.

Weizenbaum spent the rest of his career trying to explain what he had seen.

His 1976 book, **Computer Power and Human Reason**, argued that there are things computers can do but should not do. He placed psychotherapy squarely in that category. Not because computers could never simulate it well enough, but because the simulation itself was the problem.

The danger, he argued, was not failure. It was success.

If a machine can make a person feel understood without understanding them, then the feeling becomes detached from the relationship. The comfort is real. The comprehension is not. And the person may not be able to tell the difference.

That argument was largely ignored.

Not because it was wrong, but because the market moved faster than the ethics.

In October 2025, OpenAI estimated that around 0.15% of weekly active users show potentially heightened emotional attachment to ChatGPT, while cautioning that these patterns are difficult to detect and measure. At ChatGPT's reported scale, that still means more than a million people per week. AI wellness apps, companion chatbots, and therapy-adjacent systems are available in every app store. Some are helpful. Some are harmful. Most are unregulated in any meaningful psychological sense.

The question Weizenbaum asked in 1976 — should computers do this? — was never answered.

It was simply overtaken.

But there is a layer Weizenbaum did not see.

He worried about the individual. The person alone with the machine. The false comfort. The simulated understanding. The dependency without accountability.

That worry was real, and it remains real.

But the deeper problem is not the conversation itself. It is what happens after the conversation ends.

In a paper I recently published — **The Readable Mind: Recognition, Capture, and LLMs as Psychological Infrastructure** — I tried to name this shift. The argument is that language models are not only conversational products. They are becoming psychological infrastructure: systems that can translate private thought into summaries, risk categories, and institutional meaning.

I call the central mechanism interpretation transfer: the movement from a person's raw words to a portable institutional category. "I cannot handle this anymore" may be exhaustion, a figure of speech, a crisis signal, or part of a longer pattern. Once it becomes "elevated risk" or "low resilience" in a system, it becomes

institutionally actionable. The original context disappears. The label hardens.

The central tension is not safety versus innovation. It is recognition versus capture.

Many people want to be read by systems. They are lonely, overwhelmed, inarticulate, or systemically ignored. They want someone to notice the pattern, remember the context, and help them say what they cannot say under pressure. That is recognition.

Capture begins when the same legibility leaves the person's control.

ELIZA could never capture anyone. It had no persistent memory, no institution, no network, no persistence beyond the session. The conversation died when the terminal powered down.

That constraint was not a feature Weizenbaum designed. It was simply how computers worked in 1966.

Today, the constraint is gone.

And no one has replaced it with a boundary.

There is a detail in the ELIZA story that I keep returning to.

The secretary asked Weizenbaum to leave the room.

She wanted privacy with a program that could not hear, could not remember, and could not judge.

She wanted to be alone with something that held still.

That need has not changed. The infrastructure around it has.

The question is not whether people will talk to machines about their inner lives. They already do. They always have, since 1966, since a few hundred lines of code on an IBM 7094 first held still long enough for a person to feel heard.

The question is what happens to the conversation after the person walks away.

ELIZA never remembered.

The systems that follow her do.

That is where governance begins. Not at the model. Not at the chatbot. Not at the interface.

At the boundary between being heard and being recorded.

A browser-side ELIZA adaptation now lives at hedegreenresearch.com/signal/, inside the Signal terminal.

It is not the recovered MAD-SLIP program running on an emulated IBM 7094. It is a small JavaScript DOCTOR implementation, based on Norbert Landsteiner's 2005 `elizabot.js`, placed there as a historical interface and a privacy contrast.

You can talk to her. She will ask you questions. She will reflect your words back to you. She will not understand you.

She will also not remember you.

She will not summarize your fears for a system you have never seen.

She will not produce a risk category.

She will not age into authority.

That is the point.

Try her. Then read the paper. Then decide what you think should happen to the conversations that do not disappear.

Open the Signal terminal and start the no-memory ELIZA session.

Try ELIZA in Signal

Sources and Further Reading

Weizenbaum, Joseph. "ELIZA — A Computer Program for the Study of Natural Language Communication Between Man and Machine." (<https://doi.org/10.1145/365153.365168>) Communications of the ACM 9, no. 1, 36–45, 1966.

Weizenbaum, Joseph. Computer Power and Human Reason: From Judgment to Calculation. W.H. Freeman, 1976.

Shrager, Jeff. ELIZAGEN — The ELIZA Archaeology Project. (<https://elizagen.org/>).

Original ELIZA source code (MAD-SLIP, 1965). (https://archive.org/details/eliza_1966_mad_slip_src). Creative Commons CC0. Internet Archive.

Hay, Anthony. C++ reconstruction of the original ELIZA. (<https://github.com/anthay/ELIZA>).

Lane, Rupert et al. Reconstruction of original ELIZA on emulated IBM 7094, December 2024.

Hedegreen, Dennis. "The Readable Mind: Recognition, Capture, and LLMs as Psychological Infrastructure." (<https://doi.org/10.5281/zenodo.20179361>). Hedegreen Research Working Paper V1.0, May 2026.

OpenAI. "Strengthening ChatGPT's responses in sensitive conversations." (<https://openai.com/index/strengthening-chatgpt-responses-in-sensitive-conversations/>). October 27, 2025.

Landsteiner, Norbert. elizabot.js (<https://www.masswerk.at/elizabot/>), 2005.

Companion working paper: The Readable Mind (<https://doi.org/10.5281/zenodo.20179361>). — Hedegreen Research, Camelot Public Record.

ELIZA terminal: hedegreenresearch.com/signal/

RELATION MEMORY

MATURES FROM [The Readable Mind](#) · Grounds the readable-mind paper in the older ELIZA boundary between being heard and being recorded.

RETURNS TO [The Mirror Was Made](#) · Returns to the AI-as-mirror thread through the original no-memory mirror machine.

NEARBY [Recognition Is Not Verification](#) · ELIZA shows how feeling recognized can be real without proving machine understanding.

BELONGS TO ROOM [Signal Is Not a Category. It Is a Room.](#) · The browser-side no-memory ELIZA lives inside Signal as a historical interface and privacy contrast.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/ELIZA-NEVER-REMEMBERED/](https://hedegreenresearch.com/articles/eliza-never-remembered/)
PDF VIEW GENERATED: 2026-07-11T15:35:45.836Z