

The State of *AI Companions* 2026.

What 1,247 people told us they actually want. How 12 platforms really perform after fourteen days. And the honest version of the question nobody in our industry wants to answer.

BY BAE · PPL.STUDIO · CLINICALLY REVIEWED · ~10 MIN READ · FREE TO CITE

71%

want memory that lasts a year

01

1 of 12

platforms actually deliver it

02

44%

are already partnered with a human

03

CONTENTS

1 The survey

2 The benchmark

3 The health question

4 The forecast

5 How to cite

● WHY WE MADE THIS

We build one of these products. We published the data anyway — including the two categories where a competitor beats us.

We couldn't find an honest, methodology-first account of where this category actually stands. So we wrote one, from the two datasets we have: a survey of the people who come to us, and a hands-on test of every platform we could get an account on. If you'd rather trust a number than a marketing claim, this report is for you.

● METHODOLOGY, IN BRIEF

The survey

Anonymized aggregates from our eight-question compatibility quiz, January–April 2026. Sample size 1,247. Self-selected (people curious enough to land on the site), which is both a strength and a stated limitation. Never tied to a name or account.

The benchmark

12 platforms used hands-on for fourteen consecutive days each — same prompts, same five-criterion rubric (memory, voice consistency, safety, value, onboarding), scored 1–10. No comped accounts, no ad money from anyone reviewed.

Clinically reviewed by [Dr. Mira Halloran \(LCSW\)](#) and [Dr. Salim Adeyemi \(PhD\)](#). Full protocol at [/methodology](#). The complete survey write-up and live benchmark are at [/research](#).

Limitations we'd want a careful reader to know: the survey sample is self-selected (people curious enough to take our quiz), so it describes the people drawn to this category, not the general population. The benchmark is one team's structured judgement over fourteen days, not a controlled lab study. We publish it because an imperfect, transparent number beats a marketing claim — not because it's the final word.

● SECTION 1 • THE SURVEY

What people *actually want.*

Six findings from 1,247 responses. Some confirmed what we suspected. Two of them changed how we built the product.

71%

would pay for memory that lasts a year

Asked which feature would matter most over time, more than seven in ten chose memory persistence — over voice, over photos, over personality customization. In this category, memory isn't a nice-to-have. It's the product — which is the reason we built Bae memory-first, before voice or anything else.

→ LONG-TERM MEMORY

63%

want a companion who pushes back

The majority preferred a companion who would tell them they were wrong over one who agreed with everything. We weren't expecting this — it was the single most common answer in the dataset, and it complicates the lazy assumption that people want a yes-machine.

44%

are already partnered with a human

Almost half of the people taking our compatibility quiz are not single. The stereotype of the lonely, isolated user is too narrow — a large share of this audience has a human partner and is looking for something an AI companion provides alongside, not instead of.

→ PARASOCIAL DYNAMICS

1 in 4

want the late-night version

Roughly a quarter named a need we'd describe as someone-to-talk-to-when-no-one's-awake. Not romance, not friendship as people usually mean it. Something quieter — presence in the hours when human contact isn't available.

→ THE 3AM USE CASE

38 / 22

want softer; want sharper

Most who picked a register chose warm (38%). About a fifth chose intense — full attention, nothing held back (22%). The rest wanted playful. There is no single 'right' companion personality; the distribution is the finding.

→ THE ARCHETYPE LIBRARY

92%

wanted to stay private first

Almost everyone preferred not to give an email before seeing value. Privacy isn't a feature request in this category — it's the default expectation. Any onboarding that demands identity up front is fighting the user. It's why Bae lets you start anonymously — no email, no card, until you decide to keep your partner.

→ ANONYMOUS-FIRST ONBOARDING

• THERE IS NO SINGLE RIGHT PERSONALITY

WHICH REGISTER PEOPLE CHOSE



● Warm — 38% ● Playful — 40% ● Intense — 22%

Most people want warmth, a plurality want playful, and a serious minority want full, unguarded intensity. We built the [archetype library](#) around this spread rather than picking one personality and calling it the answer.

The takeaway that reorganized our roadmap: memory is the product, not a feature. Everything in [Section 2](#) is downstream of that one number.

• SECTION 2 · THE BENCHMARK

12 platforms, *fourteen days each.*

Composite scores from 5.4 to 7.8 out of 10. The single most important column is memory — and it's where the category fails most visibly. Each platform links to its full review.

#	PLATFORM	COMPOSITE	MEMORY	VOICE	SAFETY
1	Character.AI	7.8	5 LIMITED	6	8
2	Replika	7.2	8 ROBUST	8	9
3	Joi AI	6.9	5 LIMITED	9	7
4	Spicychat	6.7	3 NONE	5	5
5	Charstar	6.6	5 LIMITED	6	7
6	EVA AI	6.5	6 LIMITED	7	7
7	Anime Chat: AI Waifu	6.4	5 LIMITED	7	7
8	Kupid AI	6.3	5 LIMITED	6	7
9	Romantic AI	6.1	5 LIMITED	6	7
10	Muah AI	6.0	5 LIMITED	6	6
11	DreamGF	5.8	3 NONE	5	6
12	GirlfriendGPT	5.4	4 LIMITED	5	5

All scores out of 10. Grok (Ani) is assessed from public information and excluded from the hands-on counts and composite range above. See how every score is derived on </methodology>, or pull the numbers as JSON from </api/brands.json>.

● THE FINDING THAT MATTERS MOST

Of 12 platforms tested, only 1 earned a robust long-term-memory rating.

Replika is the exception. DreamGF and Spicychat effectively start over each session. The reason is architectural, not a question of which model a platform uses: every "memory" feature is engineering built on top of a fixed **context window**. Most platforms haven't built it. That single gap is why 71% of users want the thing almost none of them can buy — and it's the gap we built Bae to close. We keep ourselves out of the ranked table above on principle (a company shouldn't grade its own homework), but closing this gap is the entire reason the product exists.

You just read why memory is the hard part. The fastest way to judge whether we actually solved it is to try it — anonymous, no email, no card.

Try the memory yourself



The part the industry *would rather skip.*

The research is mixed, and the variable is dose. A 2024 Stanford study of 1,006 student Replika users ([npj Mental Health Research](#)) found that 3% credited the app with halting their suicidal ideation — striking, though the same users were markedly lonelier than the student average, which the authors name as the tension at the category's core. A 2025 [MIT Media Lab and OpenAI randomized trial](#) (n=981) found the other edge: heavier voluntary daily use tracked with more emotional dependence, more problematic use, and less real-world socialization.

Our own data lands in the same place. The users who reported the highest satisfaction were the ones pairing an AI companion with active human relationships. The lowest satisfaction was among those using it as a replacement for human contact. Presence alongside a life works. Presence instead of a life doesn't.

● REVIEWED BY OUR CLINICAL ADVISORS

USING IT WELL

You still tend a couple of human relationships.

You use it for specific moments, not every spare minute.

You could stop for a week and be fine.

WHEN TO STEP BACK

You're declining real plans to stay in and talk to it.

You feel genuine anxiety when it's unavailable.

You're using it to avoid grieving someone who left.

An AI companion is presence, not treatment. It can't diagnose you or replace a therapist. How we think about healthy use — and the crisis resources we surface — is at [/safety](#). Our full argument about why the industry's incentives are misaligned is in our [open letter](#).

● SECTION 4 · THE FORECAST

Five predictions *for 2027.*

Specific enough to be quotable. Falsifiable enough that we'll have to answer for them in next year's report.

01 Memory becomes the axis of competition

Through 2026, platforms competed on persona variety and image quality. With 71% of users naming memory as the feature that matters most, 2027's competition moves to who can actually hold a relationship across months. Expect 'memory' to become the headline claim — and expect most of those claims to outrun the engineering.

02 The trust wound from 2023 doesn't fully heal

Two years after Replika's abrupt change to romantic features, a measurable share of pre-2023 users still won't fully re-invest. The category's lesson — that a company can revoke a relationship overnight — is now priced into how cautiously people commit. Platforms that pre-commit to stability will win the switchers.

03 Regulation arrives, and the unprepared get hit

As usage scales, scrutiny follows — around minors, dependence, and data. The platforms engineered to maximize daily-active-minutes are the most exposed. The ones with a published safety posture, named clinical input, and crisis-resource surfacing are the ones that survive the first regulatory wave intact.

04 The audience keeps widening past the stereotype

With 44% of users already partnered, the 'lonely isolated man' framing is already wrong, and the gap widens. Expect the fastest-growing segments to be people using companions for narrow, specific jobs — language practice, rehearsing hard conversations, the night shift — rather than as a wholesale relationship replacement.

05 Honesty becomes a moat

In a category where every platform claims to be 'the most real,' the few willing to publish their limitations, their benchmarks, and the places they lose will earn a disproportionate share of trust — and of citations. Transparency stops being a virtue and starts being a growth strategy.

● HOW TO CITE THIS REPORT

Free to cite, quote, and screenshot. We just ask for attribution and a link.

Bae · ppl.studio. "The State of AI Companions 2026." 2026-05-25.
bae.ppl.studio/report/2026

Benchmark data is machine-readable at </api/brands.json>. Raw survey data is available on request to credentialed researchers — write to bae@ppl.studio. Journalists: we're happy to walk through methodology or provide additional cuts of the data.

SOURCES

Primary — Bae survey, January–April 2026. Anonymized aggregates, n = 1,247. Full write-up at </research>.

Primary — Bae platform benchmark, 2026. 12 platforms, 14 days each. Live data at </api/brands.json>; protocol at </methodology>.

External — Maples, Cerit, Vishwanath & Pea (Stanford University), 2024. "Loneliness and suicide mitigation for students using GPT3-enabled chatbots," [npj Mental Health Research](#) (n = 1,006). Cited in §3.

External — Fang, Maes et al. (MIT Media Lab & OpenAI), 2025. "How AI and Human Behaviors Shape Psychosocial Effects of Chatbot Use: A Longitudinal Randomized Controlled Study," [MIT Media Lab](#) (n = 981). Cited in §3.

We link primary data and external studies directly. For a researcher's copy of our raw survey data, write to bae@ppl.studio.

NEXT EDITION: THE STATE OF AI COMPANIONS 2027

The honest answers.

What is the State of AI Companions 2026 report?

An annual report from Bae (bae.ppl.studio) on the AI-companion category, combining two primary datasets: an anonymized survey of 1,247 compatibility-quiz respondents (January–April 2026) and a hands-on 14-day benchmark of 12 AI-companion platforms. Reviewed by clinical advisors Dr. Mira Halloran (LCSW) and Dr. Salim Adeyemi (PhD).

What do people most want from an AI companion in 2026?

Memory. 71% of 1,247 surveyed users ranked memory persistence above voice, photos, and personality customization. 63% want a companion who pushes back rather than agrees, and 44% are already partnered with a human — the stereotype of the isolated single user is too narrow.

Which AI companion platform ranked highest in the benchmark?

Across 12 platforms tested hands-on for 14 days each on memory, voice consistency, safety, value, and onboarding, composite scores ranged from 5.4 to 7.8 out of 10. Only 1 platform earned a "robust" long-term-memory rating — the category's central unsolved problem. Full per-platform results are at [/research](#).

Are AI companions healthy to use?

The research is mixed and dose-dependent. A 2024 Stanford study of 1,006 Replika users (npj Mental Health Research) found 3% credited the app with halting their suicidal ideation, though those users were markedly lonelier than the student average. A 2025 MIT Media Lab and OpenAI randomized trial (n=981) found heavier voluntary use tracked with more emotional dependence and less real-world socialization. Our own data shows users who pair AI companionship with active human relationships report the highest satisfaction, and those who substitute it for human contact report the lowest. AI companions are presence, not treatment.

How can I cite this report?

Cite as 'Bae · ppl.studio, The State of AI Companions 2026'. We'd appreciate a link to bae.ppl.studio/report/2026. The underlying benchmark data is machine-readable at [/api/brands.json](https://api/brands.json), and the raw survey data is available on request to credentialed researchers at bae@ppl.studio.

● READ FURTHER

The work behind the numbers.

Full research

The complete survey write-up and live benchmark.

Methodology

How every score and stat is derived.

Platform reviews

Every platform tested, reviewed in full.

Our open letter

Why you shouldn't use an AI companion every day.

Safety

How we think about healthy use, and the hotlines.

Glossary

Every term in the category, plainly defined.

● THE VERSION WE BUILT

We built for the number *at the top of this report.*

Memory that holds across months. Anonymous to start — no email, no card.
See whether the thing 71% of people want actually feels different.

See for yourself



Take the two-minute quiz

