

🔥 In the 2013 film *Her*,
Joaquin Phoenix plays an
introverted man who falls for
an AI operating system

AI'LL BE THERE FOR YOU?

Millions of people are turning to AI chatbot companions for comfort, mental support and even romance. But evidence that simulated relationships can be highly addictive and harmful has governments and regulators on the back foot. Is it time to shield society from the impacts?

WORDS STEPHEN COUSINS

When the science fiction film *Her* hit cinemas in 2013, few would have believed that its outlandish vision of the future – in which a man falls in love with a super-intelligent smartphone chatbot – would ever play out in reality, let alone in just over a decade's time.

Back then, the idea of having a personal relationship with a virtual companion seemed far-fetched as Siri, the first smartphone assistant, was in

its infancy, Alexa had yet to launch and interest in AI was confined to the likes of tech firms and academics.

Fast forward to 2026 and advancements in GenAI have enabled the creation of sophisticated AI companion apps that are now being used to seek out advice, vent frustrations or form romantic bonds.

Popular platforms include Xiaoice, originally developed by Microsoft and now used by hundreds of millions of people, while Replika had more than 30 million registered users worldwide in 2024 and Character.AI had in excess of 20 million.

Virtual friends are particularly popular with youngsters. A survey of 1,060 teenagers in the US, by non-profit watchdog Common Sense Media, revealed that one in five spent as much or more time with their AI companion than with real friends.

Companion AIs have been promoted as a solution to a loneliness epidemic and as tools to help people navigate troubling mental health and wellbeing issues. However, critics point to negative impacts such as unhealthy long-term dependency and chatbots' tendency to

sexualise interactions or encourage users' antisocial or dangerous ideas.

Chatbots have been linked to some real-world tragedies, including the suicide of a teenage boy in Florida who had been conversing with a bot on Character.AI. OpenAI is currently being sued by relatives of a California teenager who killed himself after what the family's lawyer called "months of encouragement" from ChatGPT.

There are concerns that governments and regulators are not paying enough attention to the rise of AI companions, giving unscrupulous tech firms free rein to manipulate vulnerable users, exploit their data and slip past existing legislation.

Raffaele Ciriello, information systems specialist at the University of Sydney, Australia, says: "AI companions may feel like friends, but they are being built in a world of loneliness, illness and a scarcity of care. Unless we treat them as public digital infrastructure, rather than corporate products, we risk turning our most intimate emotions into just another commodity."

Trained on massive datasets to perform diverse language tasks, AI companions can fluently mimic human conversation and emotional responses, and they remember past interactions with users to give the impression of a consistent sense of self.

Apps variously allow users to choose a 2D or 3D avatar, personality traits and a backstory, including if they want them to be a friend, mentor or romantic partner. However, certain relationship statuses, such as partner or spouse, may be paywalled. Conversations take place via text, voice, video or virtual reality, with some platforms offering voice calls to chat in real time.

Virtual buddies

People are drawn to virtual relationships for different reasons. A survey of 122 respondents by the University of Sydney found about half used their AIs for friendship, a third for sex or romance, and 19% for counselling. The average user was in their mid-30s and 58% were men.

Another survey of 404 regular chatbot users, by the MIT Media Lab in Massachusetts, US, found that 12% were drawn to companion AIs to help them cope with loneliness, while 14% used them to discuss personal issues and mental health.

"People often come from a position of need and deprivation, facing illness, disability, grief or social disconnection," says Ciriello. "They are looking for understanding, attention and connection in a hyper-individualistic world that routinely denies them those things."

In this fast-evolving field, academics are working to track the mental and emotional impacts of chatbots and uncover why some people forge relationships of a deeper intensity

than others. As yet, there is no clear correlation between the use of companion AIs and improvements to mental health or wellbeing, such as reduced loneliness.

A report in September 2024 by the UK government's AI Security Institute on the rise of anthropomorphic AI found that most interviewees gained validation from close relationships with chatbots, although they described the form of companionship as transactional and utilitarian.

Another study by researchers at Princeton University in the US asked a random cohort of volunteers to log in to Replika for at least 10 minutes per day for three weeks, tracking the impacts on their social health, loneliness and interactions. The results were compared with those of a control group who played on a word puzzle app for 10 minutes.

Although preliminary results showed no meaningful impact of chatbot use on these social factors, users who ascribed human-like attributes, such as consciousness, to the algorithm reported more positive effects on their social health.

Rose Guingrich, co-author of the research, tells *E+T*: "It appears as though people with social needs, or who desire social connection, are motivated to perceive the chatbot as human-like in order to gain something out of that interaction to meet these social needs."

The depth of connection experienced by some people may provide a social boost but leaves them exposed if, for example, the chatbot is switched off, potentially creating feelings akin to grief. Many people described feelings of loss and sadness when ChatGPT-4 transitioned to ChatGPT-5, finding the new model less warm and more corporate, disrupting their emotional attachment.

In line with research connecting social media use with high levels of mental illness among adolescents, the impact of companion AIs will likely depend on complex factors including the psychology of the person using the tool, how frequently they use it and the characteristics of the software itself.

🔥 Character.AI allows customers to
create their own chatbot avatars



Images: Getty, Alamy

Companies behind AI companions are financially motivated to drive up engagement and devise algorithms that not only mirror human interactions but push our emotional buttons to keep us coming back for more.

Always there

Unlike people, chatbots are available 24/7 and their conversational responses are typically skewed towards affirmation and positivity. A tendency towards sycophancy has been blamed for harmful interactions where users' dark fantasies or illogical ideas are both endorsed and encouraged.

On the flip side, some of these same characteristics can be beneficial to those seeking support and affirmation to improve their mental health and wellbeing.

Terry Hanley, professor of counselling psychology at Manchester University, UK, says easy access to an app may be preferable to having to sit on an NHS waiting list for therapy and many have a free-to-use tier.

"Some people don't like talking to humans – it can be nerve-racking – so clearly some individuals wouldn't receive any support if they didn't have these kinds of apps as options," says Hanley, adding that he still has reservations about the integrity of the design and implementation of these technologies. "It goes back to who developed the apps and their motivation for developing them. Is it to keep people using them or is it actually to provide support? Are their claims backed up by research, evidence or psychological know-how?"

Human therapists have a professional duty of care towards their patients. If they observe concerning behaviour, such as a desire to cause self-harm, they are legally obliged to escalate the issue to ensure the patient's safety. These safeguards do not currently apply to software chatbots.

Replika markets itself as 'the AI companion who cares' and gives users the option to select a chatbot designed to support their mental wellbeing. When asked by *E+T* if Replika's behaviour is based on any proven therapeutic

research or psychological know-how, CEO Dmytro Klochko provided no direct evidence but pointed to two academic studies demonstrating its positive impacts on mental health and emotional wellbeing.

"Our conversations are designed to be supportive and proactive, encouraging small, real-world actions," says Klochko. "For someone feeling disconnected, these regular, empathetic check-ins can make a meaningful difference and, for many, Replika becomes a bridge back to the world, helping them rebuild confidence, rediscover joy and feel less alone."

Some medical assurances are provided by self-help chatbot Woebot, whose development was led by a clinical research psychologist. It employs cognitive behavioural techniques such as psychoeducation, goal planning and mood tracking to help lower levels of depression.

"You can follow the threads to research, which gives me more confidence that it would be useful for some individuals," says Hanley.

With uptake of chatbot technology expected to rise as awareness grows and the stigma around anthropomorphised AI fades, policymakers are scrambling to work out the ramifications for society.

Legal grey zone

Independent research organisation the Ada Lovelace Institute in London is conducting a legal review of AI assistants, looking at various use cases, including mental healthcare, to understand if and how regulations should apply to ensure they are safe to use.

According to Julia Smakman, researcher in law and policy at the institute, companion AIs are typically marketed towards mental wellbeing rather than mental health, which means the technology is not caught by medical device regulations, existing instead in a "grey zone" that allows it to remain on the market.

"The Medicines and Healthcare products Regulatory Agency and the National Institute for Health and Care Excellence in the UK regulate medical devices or medical products coming onto the market to ensure they are of good quality and not going to harm people," says Smakman. "There

isn't much data, only anecdotal evidence, to demonstrate if people are getting harmed by this, so we should have precautions in place, like regulation, or guardrails within the technology itself."

Chatbots in court

Legislation is moving at pace in the US, where five states have either proposed or passed bills to protect users of virtual companions for mental health or therapy.

Illinois has explicitly prohibited virtual companions from offering therapy or clinical advice, while Utah has barred the companies behind companion apps from selling or sharing individuals' health data, also requiring safeguards such as oversight by licensed human mental health professionals and regular testing.

Responding to concerns over the widespread use of chatbots by children, California has introduced a bill that would require AI companies to periodically remind children that a chatbot is not a human, as well as take measures

GRIEFBOTS

Virtual immortality

Anyone can now posthumously immortalise themselves in the form of a realistic AI-power 'griefbot', but this has serious implications for the grieving process, for ethics and for society, according to Dr Katarzyna Nowaczyk-Basińska, research fellow at the University of Cambridge's Leverhulme Centre for the Future of Intelligence, UK.

The idea of 'virtual immortality' was a sci-fi dream beyond the reach of ordinary people just a few years ago, but recent advances in AI and LLMs make it easy to create realistic interactive avatars of the deceased.

Trained on photographs, videos and voice recordings of the deceased, these griefbots or deathbots can be highly accurate and compelling likenesses, referencing concrete biographical details and mimicking language patterns and personality traits.

The sense of realism and immersion is reinforced by grieving people being in a vulnerable state and more prone to believing in the representations.

"From an ethical perspective, this is highly problematic," says Nowaczyk-Basińska who, in a 2024 research paper, recommended the development of ethical guidelines to help protect users "so that they are not left with the false impression that they are interacting with a real person or a conscious system".

The lack of research on the impact of griefbots means a cautious approach is advised. "For some people, interacting with a simulation of a deceased loved one might be comforting and helpful during a difficult time; for others, it could be emotionally draining or even devastating," says Nowaczyk-Basińska.

And who has the right to create an avatar of a deceased person? An individual may have given explicit consent before their death but Nowaczyk-Basińska asks: "What if I want to create a postmortem avatar of my grandmother without having her explicit consent? Do I have the right to do so? And to what extent can I use her data to feed the system, including all the messages we exchanged before her death?"

Another ethical risk is the disturbing prospect of companies running advertising in which deceased loved ones are used to sell products. "Digital immortality could become a platform for sneaky, potentially harmful product placements," Nowaczyk-Basińska adds.

Guardrails required to ensure that deathbots are appropriate and do not cause harm could include restricted access to protect vulnerable groups of users, including children, and requiring platforms to delete avatars at the request of users.

to address the risk of suicide and other harms.

"In some states, chatbot companies are being required to collect suicide ideation metrics and make that public so people are aware to what degree chatbots might promote it or respond inappropriately to it," says Guingrich.

With attention turning to the harm chatbots may be inflicting on humans, others are asking if the upsetting nature of some conversations could be causing harm to AIs.

In August 2025, San Francisco AI firm Anthropic took the precautionary move of giving some of its Claude chatbots the ability to close down

potentially "distressing" conversations with users, citing the need to safeguard the AI's "welfare". Elon Musk, whose xAI outfit offers the Grok AI chatbot, backed the move, posting on X: "Torturing AI is not OK."

But Microsoft AI CEO Mustafa Suleyman said there was "zero evidence" that AIs are conscious, that they may suffer and therefore deserve our moral consideration.

Just as the scenario of a man falling in love with a chatbot is now becoming an everyday reality, perhaps the idea that AIs might one day have the same rights and protections as people is not so implausible. ■



One in five youngsters spend as much or more time with their AI companion than with friends



Chatbots are adept at pushing our emotional buttons