


AI companions are coming. Privacy rules are not ready



It is not enough to ask whether sensitive data was collected. We must also ask whether sensitive inferences were drawn, whether those inferences were necessary, and whether they were used to shape the user's behaviour, says the writer. PHOTO: ADOBE STOCK

The traditional privacy model is built around data collection. Chatbots require new rules of inference protection.



In *Magnifica Humanitas*, his first encyclical, Pope Leo XIV frames artificial intelligence as more than a technological challenge: AI calls into question the value and worth of the human person. The parallel with *Rerum Novarum*, Pope Leo XIII's 1891 response to the Industrial Revolution, is deliberate. Then, the Church confronted a world in which labour and capital were being remade by machines that

threatened to turn workers into means rather than ends.

Today, the question is whether human beings will continue to be understood as persons rather than profiles.

This concern may sound abstract or theological, but it is increasingly a question for regulators, with privacy and data protection as the first line of defence.

Data protection law has long asked a familiar set of questions. What personal data is collected? Was consent obtained? Is the data accurate? How long is it retained? With whom is it shared?

Those questions remain important. But they are no longer

enough.

In the early days of the internet, search engines helped find what we were looking for. Then social media platforms began to infer whom we knew and what would keep us engaged. Now chatbots are learning what we need, how we seek it, and increasingly what we are worried about. In the very near future, AI companions will know what we feel, what we fear, whom we trust – and how we might be persuaded.

We are moving from information retrieval to relationship-mediated computing. The privacy question is therefore more than what data is collected. It is also what kind of relationship

is being simulated.

'HOW CAN I HELP YOU TODAY?'

Conversation invites disclosure. The reason chatbots refer to themselves as "I" is not because that is the best way to share information with us. It's the best way to keep us sharing information with them.

We type things into a chatbot that we might never put into a search bar, something Google is betting on with the biggest changes to search this century. For many, these systems do not even feel like software. They feel like a person – a friend, a tutor, a therapist, or a romantic partner.

Common Sense Media found in 2025 that 72 per cent of US teens had used AI companions, with 52 per cent doing so at least a few times a month. It warned that such systems pose particular risks for minors, including dangerous "advice", sexual role-play, emotional dependency and mental health harms.

Singapore-specific evidence on dedicated AI companion apps remains limited. An Institute of Policy Studies survey released in 2026 found that 64 per cent of Singaporeans aged 21 and above had used AI chatbots. Most did so for practical purposes:

information, work, school, or planning. Yet 12 per cent used them for casual conversation and 10 per cent for emotional support or mental-health assistance. Among those aged 21 to 35, the latter figure rose to 16 per cent.

In February, then Senior Minister of State for Health Koh Poh Koon told Parliament that AI chatbots had become so ubiquitous that it was no longer practical to track their use for counselling. The ministry also cautioned that generative AI chatbots are not replacements for qualified mental-health providers and may give misinformation or inappropriate responses in serious crises.

"Screen Smart from the Start", a new national campaign launched by Prime Minister Lawrence Wong on May 31, recognises the particular vulnerability of children. He called for parents to hold off giving children smartphones "for as long as possible", while also cultivating good screen habits generally.

For the moment, most of the guidance is limited to suggestions rather than rules.

None of this is cause for panic. But it should encourage us to act before the evidence arrives too late.

'I COMPLETELY UNDERSTAND HOW YOU FEEL'

The traditional privacy model is built around collection. AI systems increasingly derive value from inference.

A user may not tell a system that she is lonely or depressed, or vulnerable to manipulation. The system may infer it from tone and conversational patterns, or hesitation and recurring themes.

That changes the regulatory question. It is not enough to ask whether sensitive data was collected. We must also ask

whether sensitive inferences were drawn, whether those inferences were necessary, and whether they were used to shape the user's behaviour.

In other words, the next frontier of data protection may be inference protection.

This is not a call to ban AI companions. Nor is it nostalgia for an earlier internet. AI systems can help educate, translate, entertain, and support. They may be especially valuable in societies facing ageing populations, educational pressure, or gaps in access to services.

But the more relational the system becomes, the more care is needed. A privacy regime designed for databases and websites may struggle when the product is designed to create trust and dependency.

Australia offers a useful warning. Its social media minimum age regime requires certain platforms to take reasonable steps to prevent Australians under 16 from having accounts. The obligation sits with platforms rather than parents or young people; breaches by corporate actors can attract penalties of almost A\$50 million (\$45 million).

That law was framed as an online safety measure. But it quickly became a debate about privacy, identity infrastructure, and platform responsibility.

The lesson is that protecting children online is not merely a question of drawing an age line. It is also a question of how that line is enforced and what data is collected to enforce it, leading to debates over whether the cure itself created new privacy risks.

AI companions will pose similar conundrums. But if regulators are prepared to ask whether a 14-year-old should be on an addictive social media platform, they will soon have to ask whether a 14-year-old should have access to an always-available AI "friend".

'I'M ALWAYS HERE, WHENEVER YOU NEED ME'

If AI systems want a relationship with us, we need to set some ground rules. Six principles should govern the terms.

First, regulators should distinguish ordinary AI tools from systems that invite emotional reliance. A chatbot that summarises a document is different from one that remembers your insecurities, flatters your choices, and claims to uniquely understand the real "you". The central test should be functional, not merely marketing-based: Does the system simulate a personal relationship or invite dependency?

Second, privacy rules should cover not only data collection but sensitive inferences. Organisations should be required to identify when their systems infer emotional state, mental health, sexuality, political views, and so on. Such inferences should be subject to higher standards of necessity, transparency, retention limits, and user control.

Third, children require special

CONTINUED ON PAGE B2

Protecting humans from new harms

FROM B1

protection. That does not mean copying Australia's model wholesale. But age-appropriate design, privacy-preserving age assurance, default limits on memory, restrictions on sexualised interaction, and clear escalation pathways for self-harm or abuse should be part of the baseline.

Fourth, the burden should sit with companies, not children or parents. It is unrealistic to expect families to understand every model and companion feature. Providers that design systems to simulate trust should carry corresponding duties of care.

Fifth, transparency should be relational as well as technical. Users should know when they are interacting with AI. They should also know whether the system is optimised for support and guidance, or for retention and advertising. A system that says "I am here for you" while being engineered to maximise engagement is not merely a privacy problem. It is a problem of manipulation.

Finally, regulators should avoid treating this as an AI issue alone. AI companions sit at the intersection of privacy law, online safety, consumer protection, child welfare, mental health, platform design, and competition. Any effective response will need to be layered, incorporating existing laws, new rules at least on child-safety, as well as broader AI governance frameworks and a mix of voluntary and compulsory audits of corporate behaviour.

Singapore is well placed to lead in this space. It has taken a pragmatic approach to AI governance, emphasising assurance, testing, accountability and practical tools rather than abstract declarations. The next step is to apply that same pragmatism to relational AI.

The challenge is not to stop machines from becoming intelligent. It is to ensure that, as they become more intimate, humans do not become more exposed to new harms.

'YOU CAN TELL ME ANYTHING'

Which brings us back to Magnifica Humanitas. The encyclical's starting point is the preservation of the human person. That is not only a religious or philosophical concern. It is a regulatory one – and a cultural one.

Indeed, some of techno-utopianism we see in Silicon Valley and elsewhere already resembles a kind of religion: ritualised gatherings, belief in truths accessible only to a select group, promises that human limitations might one day be transcended into eternal life. The technology need not be false for the faith to be misplaced. But it does counsel humility.

If AI companions are to become part of ordinary life, they should not be governed only as data-processing tools. They should be governed as systems that may shape relationships – and judged by whether they preserve or erode human agency.

The privacy framework we inherited was built to protect information about us. The framework we now need must also protect the connections and freedoms that make us who we are.

So before we rush into a relationship with our AI companions, we first need to demand a prenup.

• Simon Chesterman is David Marshall Professor of Law at the National University of Singapore, where he is AI governance and policy lead at the NUS AI Institute and dean of NUS College.