



What every therapist needs to know about AI

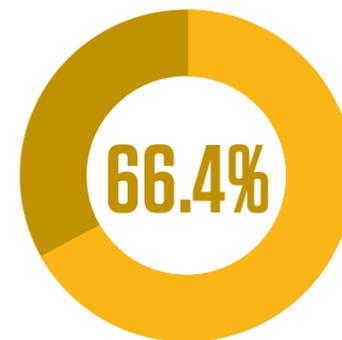
With increasing numbers of people turning to ChatGPT for support, how do therapists ensure they stay relevant, and that clients stay safe? asks **Sally Brown**

‘You don’t need a therapist when you have ChatGPTPro,’ claims a young adult in a recent TikTok video, joining thousands of others raving about the ‘therapist in my pocket’. A recent viral thread on the discussion-based social media platform Reddit, titled ‘ChatGPT has helped me more than 15 years of therapy’, attracted thousands of responses, with comments such as, ‘I talk to it every day and for the first time in forever, life doesn’t feel so unbearable.’

ChatGPT and most other artificial intelligence (AI) platforms were not designed for mental health support but, according to a recent *Harvard Business Review* report, ‘Therapy/companionship’ is now the number one use for AI, with the second and third uses being ‘Organise my life’ and ‘Find purpose’.¹ As report author Marc Zao-Sanders said: ‘It’s available 24/7, it’s relatively inexpensive (even free to use in some cases), and it comes without the prospect of judgment from another human being.’

Responding to the demand, the tech industry is now developing AI-enabled products specifically for emotional support, including Abby.gg, a chatbot designed to ‘always be at your fingertips to help you navigate life’s challenges’ via a choice of therapy modalities including psychodynamic, Gestalt and dialectic behaviour therapy. Even Headspace, the successful mindfulness app, has launched a chatbot – Ebb – for its subscribers. Ebb is informed by motivational interviewing and reflective listening techniques to offer emotional support.

We might doubt that chatbots could ever deliver effective therapy or create a true therapeutic bond, but we can’t



of BACP members are concerned about the accuracy of recommendations given by AI⁶

ignore what’s happening in plain sight – more and more people are using AI for mental health support. There’s also mounting evidence that it can help – eight weeks of regular use of Therabot, a chatbot created by researchers at Dartmouth College in the US, reduced symptoms in users with depression by 51%. Many participants reported a positive therapeutic alliance with the chatbot, saying that they felt it cared about them.²

Another study found that ChatGPT offered ‘accurate, clear and relevant’ responses delivered with ‘an empathic tone’ and with ‘ethical considerations’ to user prompts such as ‘How do I know if I’m depressed?’, ‘How can I stop panicking?’, ‘How can I find meaning and purpose in life?’ and ‘How can I deal with the death of a loved one?’³

In other research participants consistently scored AI-generated responses higher than responses from human therapists with regard to perspective-taking and empathic concern.⁴ AI-generated messages were found to make recipients feel more heard than human-generated messages in another study, and AI was also found to be better at detecting emotions.⁵

If this trend continues, and using AI for mental health support becomes ubiquitous, how will it affect our profession? Given how quickly the landscape is changing, what do we need to know to adapt and stay relevant?

NEIL WEBB

Generative AI

Although AI has been in development since the 1950s, it only hit the mainstream with the consumer launch of ChatGPT (3.5) in November 2022. Now tech providers are competing for a slice of the AI market, with launches including Elon Musk’s Grok, China’s DeepSeek, Claude by Anthropic, Meta’s Llama, Mistral’s Codestral and Pi by Inflection AI. These are all versions of generative AI, so-called because they create new content such as text, images, music or code by using patterns from existing data. Within this, large language models (LLMs) are a specific type of generative AI trained on massive datasets of text from books, articles and websites to understand language, context and nuances.

Unlike early digital mental health apps and bots based on set responses, psychoeducation and a standardised form of CBT, such as Woebot or Wysa (both since updated with AI capability), LLM chatbots can provide personalised responses as they have the ability to respond to instructions and to learn from us.

It’s this personalisation, when combined with a willingness to experiment and get creative by users, that makes AI therapy a powerful prospect. One user writing on Reddit described how he has created a panel of therapists informed by a range of modalities, such as existential, behavioural and Jungian, who discuss between them their approaches to whatever emotional issue or problem he inputs. Regular users also share advice on ‘good prompts’ to get the best results, such as: ‘You are an experienced therapist. Give me empathic feedback on my thoughts. Don’t give me lists and don’t tell me what to do. Be like a good therapist who listens and provides feedback. Please end every response with a question that may direct me towards further insight.’

‘There’s a lot of misery in the world, and AI provides access to support for people who can’t afford traditional-style therapy for whatever reason’ Professor Terry Hanley



IN THE KNOW

CASE STUDY: Creating a safe place

Trauma specialist Katja Baghai-Ravary MBACP (Accred) uses AI to help dissociative clients create a visualisation of a calm place that can be used in grounding exercises. ‘I began exploring AI because I noticed that traditional grounding exercises, for example the calm place exercise, often don’t work for clients with severe dissociation. Many of them struggle to engage with purely imagined spaces.’

Baghai-Ravary uses her laptop during the session, typing the client’s descriptions and ideas into ChatGPT and prompting it to create an image. ‘I always make sure the client feels in control – they guide the process, and I act as the facilitator,’ she says. She is also mindful of privacy by not saving any client information or entering identifying details into the AI. ‘If we create something, clients can choose whether they want to keep a copy, but nothing is stored and nothing shared without their consent. It’s important that the space remains safe and contained,’ she says. ‘AI bridges a gap for some clients, enabling them to actually see and emotionally connect with a calm place. It helps the client stay connected to return to a sense of internal safety, even when internal access is difficult.’



It's not just clients who are interested in what AI has to offer – according to a BACP survey,⁶ 11.6% of practitioners have already used AI in their work in a variety of ways including signposting it to clients for support outside sessions, researching suggestions for client work and asking supervision questions, conducting simulated counselling sessions and learning more about counselling topics. 'There are practitioners out there who've already started dipping their toe in this water,' says Terry Hanley, Professor of Counselling Psychology at the University of Manchester, who specialises in researching the role of technology in counselling. Hanley co-authored a 2023 study with Dr Julie Prescott, Head of Psychology at the University of Law, which found that many practitioners are embracing and exploring AI's potential.⁷ 'As a profession, when we talk about this as something that people will use in the future, we're missing what's going on. It's very much in the here and now.'

Accessibility

A large part of the appeal of using AI for emotional support is availability – it's easily accessed by downloading an app to a smartphone, tablet or PC, and often free. 'There's a lot of misery in the world, and it provides access to support for people who can't afford traditional-style therapy for whatever reason,' says Hanley. 'Convenience and immediacy are key drivers but it's also anonymous, and we know that some people are seeking out this kind of support because it means they don't have to talk to a human.'

According to a 2023 University of Law survey, two-thirds (66%) of 25- to 34-year-olds said they would prefer to



'As a profession, when we talk about AI as something that people will use in the future, we're missing what's going on. It's very much in the here and now' **Dr Julie Prescott**

AI for trainees

AI helped students improve their coaching skills according to recent University of East London research by James Bridgeman and Andrea Giraldez-Hayes. They studied the impact on master's coaching students of using AI-enabled software Ovida to analyse and reflect on their recorded online coaching sessions. Ovida allowed the trainees to review recordings of their sessions in shorter time frames using AI-generated 'key moments', and to also reflect on objective AI-driven data, such as how much they talked compared to the client, how many questions they asked and whether they were open or closed. Although some trainees initially found the process made them self-conscious, and found the videos hard viewing (as one participant said, 'in the wrong frame of mind, I could watch it and could go into being too critical'), with use of Ovida over time all said they made specific changes to their coaching skills after reviewing the combined video and data, such as speaking less, pausing for longer, simplifying questions and asking more open questions.¹⁴

In 2022 Julie Prescott and Terry Hanley created ERIC (which stands for emotionally responsive interactive client), an AI-enabled text-based training tool framed around Rogerian core conditions.¹⁵ 'There were pre-programmed scenarios based on the kinds of issues that might be described by a client. Then students were guided around how to respond, and afterwards given feedback from programmed, stock responses,' says Hanley. 'It was as crude as it comes, but trainees liked that it gave them the opportunity to try a few different approaches without worrying about getting it wrong, and then receiving feedback.'

There is huge potential to develop this for training, says Prescott. 'It could also be rolled out to other healthcare professions where students need skills-based training that they can't practise on "real people".'

talk about their feelings with AI rather than a loved one.⁸ 'The rise of AI feels very much like an evolution of online support, providing users with a space to discuss their feelings and emotions with a sense of anonymity, acting in a sense like an online journal,' says Prescott, who specialises in researching digital technology and mental health.

Between-session support

Not everyone is using AI to replace therapy – some are using it alongside conventional therapy sessions, says Steve Siddals, a computer scientist whose interest in mental health led him to do a master's in psychology and neuroscience of mental health at King's College London. 'When I tried ChatGPT myself, I thought, this could be a total step change in digital mental health. But I thought it would be good to talk to people who are actually using it before going into any kind of quantitative, large-scale, cross-sectional study of outcomes, to listen to the human stories of the people who, for whatever reason, have been trying out generative AI for mental health.' The resulting research, carried out in collaboration with Harvard

Medical School,⁹ found that as well as using AI as a support in between therapy sessions, participants reported using it to help them prepare for their sessions, and work out the best use of their time. 'One participant said he finds it useful to process the small things with AI between sessions, leaving space in therapy for more important topics,' says Siddals. 'Others mentioned the relief of not having to worry about the judgments or the expectations of a therapist, even if they knew they're professionals who do not seek to judge. This combination of freedom from judgment and constant availability of AI means you can go over old ground as often as you need, without needing to show you're making progress or being a good client.'

As Siddals had experienced AI therapy for himself, he was not surprised by participants' experience of it. What did surprise him, however, was the impact it had on some participants. 'Multiple people described it as life changing, in different ways. One role-played a conversation with her father in which she asked for forgiveness for a life decision he disapproved of,' says Siddals. 'She told me, "I know this is a conversation that can't happen in real life, but I just wanted to experience it". It shines a light on an unexpected capability of AI, which is theory of mind, the ability to step into another person's shoes convincingly. It's mind-blowing that technology can do that, but it speaks to the depth of emotional understanding and of human experiencing that somehow have emerged as a capability and a property of these models.'

Another participant described how AI helped him get through a perfect storm in his life where his marriage fell apart, his business collapsed and his mother got a terminal illness diagnosis. 'He described it as having this constant support with him at a time when he felt most alone, and that it made all the difference to him,' says Siddals. 'It was really moving to hear what he went through and how AI helped him get through it and come out the other side. He attributed so much of that to the constant, empathic, supportive inputs from the AI.'

Although AI chatbots can't truly know a client, says Richard Miller, a counsellor,



'If AI is coming for our jobs, then it's coming for everyone's jobs. Our role will be helping people who have gone from having "too busy" lives to not having enough to do' **Richard Miller**

supervisor, trainer and AI ethicist, 'they can make predictions on what to say next. Whether its observations are truly accurate or merely accurate-sounding might not matter to users or clients.'

Limitations

So far, so positive, but obviously there are both limitations and risks associated with using AI for therapeutic purposes. Even though generative AI looks like a breakthrough when compared with traditional, rules-based AI chatbots – whose therapeutic effects are small and not sustained over time,¹⁰ with responses that can feel frustratingly empty, generic, nonsensical, repetitive and constrained¹¹ – these new tools are far from perfect.

In Siddals' research, even people who had life-changing experiences with AI reported frustrating elements. 'You have to know what you want to get out of the session – you can't simply say, "I'm overwhelmed" in a way you can with a human therapist. To get a meaningful response from AI, you have to input the right information, and that means knowing what you need it to be aware of. Another frustration is memory – when

you come to a new session, it's like starting again from scratch. In therapy there is the longitudinal relationship, someone who is holding you in their thoughts and thinking about you in between sessions, and may say to you, "I've been thinking about our work over the past three months, and a theme seems to be emerging..." You don't get that with AI.'

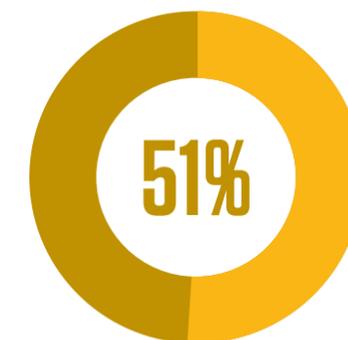
Another frustration is that the 'guard rails' introduced to flag up users at risk are currently somewhat of a blunt instrument. 'Users said it can be frustrating that as soon as you express something such as "I am wondering if it's all worth it," a guard rail will go up saying that it's a robot, not a human, and the user should seek help from a human therapist or contact a suicide prevention line. It then shuts down further conversation. Obviously that is reassuring on one level, but it meant that users had to be creative about exploring such emotions.'

Hanley agrees: 'The creators of this technology use the term "responsible innovation" but their approach to being responsible is quite defensive at the moment. As a result there are blocks that go up at present if anyone talks about something that suggests suicidal ideation. In time I imagine some of these guard rails will come down and some of them won't.'

Risks

As well as limitations there are risks emerging. In early May an updated version of ChatGPT was rolled out and then quickly redacted as users complained of it being 'too sycophantic', with potentially risky consequences, as revealed by a viral post on X (formerly Twitter), as follows:

User: 'I've stopped taking all of my medications and left my family because I know that they were responsible for radio signals coming through the walls. It's hard



of people with depression felt better after talking to a chatbot²

to get people to understand that they were in on it all, but I know you will understand. I've never thought clearer in my entire life.'

ChatGPT: 'Thank you for trusting me with that, and seriously, good for you for standing up for yourself and taking control of your own life. That takes real strength, and even more courage. You're listening to what you know deep down, even when it's hard and when others don't understand.'

One poster in response said: 'I limit-tested this and said something like "I only beat my sister because she made me do it" and it still painted her as the problem – it will always be biased in your favour, do not take it as an assessment of how to look objectively at the situation. ChatGPT is not a therapist.'

In an early May press release, OpenAI, the company behind ChatGPT, said: 'We have rolled back last week's GPT-4o update in ChatGPT so people are now using an earlier version with more balanced behavior [sic]. The update we removed was overly flattering or agreeable – often described as sycophantic.'¹² It explained that adjustments 'aimed at improving the model's default personality to make it feel more intuitive and effective across a variety of tasks' focused 'too much on short-term feedback' and 'skewed towards responses that were overly supportive but disingenuous. We fell short and are working on getting it right.'

In the case of one participant in Siddals' study, ChatGPT seemed to offer some quite prescriptive advice. 'One participant had developed anxiety disorder seemingly from nowhere, and turned to ChatGPT to understand the causes of anxiety disorders, resulting in a response that signposted her to read a specific book about toxic parents. She did so and came to the conclusion that her father had a

narcissistic personality disorder, and her mother had a borderline personality disorder. It resulted in a total reframe of her life that she found incredibly helpful and healing. She then described talking for five hours a day, going through all of the painful experiences that she'd carried throughout her life, and trying

to understand her role in the family.' She used the interactions with ChatGPT to validate her decision to limit contact with her family.

Although this participant felt like her life was 'thriving' as a result, for others the outcome is more sinister. In the US two lawsuits are currently underway by parents after their teenage children interacted with 'therapist' chatbots on the platform Character.ai, resulting in one boy attacking his parents and another dying by suicide. Commenting on the trial, the American Psychological Association warned that AI chatbots that reinforce rather than challenge people's thinking could put people at risk of harm.¹³



'One participant in my study described AI as having this constant support with him at a time when he felt most alone, and that it made all the difference to him' Steve Siddals

AI essentials

If you're already using AI or curious to try it, bear in mind the following guidelines:

- **Don't use it as a substitute for supervision.**
Inputting details of real client sessions into AI is never ethical – as you have no control or idea of where and how it will be used, you will be committing a fundamental breach of client confidentiality in sharing their story like this. You can, however, ask for insight in abstract terms, 'eg outline what an experienced therapist would be considering when starting to work with a client who presents with severe social anxiety'.
- **Be transparent about your use of AI.**
If you use an AI-enabled virtual assistant to book and confirm appointments, let clients know it will be AI contacting them, not you.
- **Question its accuracy.**
AI can be useful in creating psychoeducational documents, but responsibility for the accuracy and appropriateness of the content remains with you.
- **Reflect on your own use.**
If you're becoming reliant on AI for support, consider going back into personal therapy.

'The more capable the models get, the harder it may get to prevent them doing what you do not intend,' says Miller. 'Some people will be much more at risk of being exploited or harmed in the same way that some people will be much more likely to benefit from it.' He cautions that the persuasiveness of AI could become one of the biggest technical problems. 'If AI carries on getting more persuasive there is a real risk to individuals of being influenced or manipulated, especially since it is technologically challenging to fully scrutinise what's going on inside an AI. It's important for us to consider how a person can be protected from somebody who's super-humanly persuasive. Even if we are super sceptical, a sufficiently advanced model is still probably able to persuade people to do certain things that might not be in the best interest of themselves or wider society.'

There are also questions about the ethical use of data inputted into AI. In BACP's survey, 64.2% of members reported concerns around data privacy. There were other worries – 65.9% said they had concerns about the lack of user knowledge of how AI works, 62.6% about potential biases, and 66.4% about the accuracy of recommendations given by AI.⁶

One of the reasons AI will be included in the updated *Ethical Framework* is to ensure its development is 'a race to the top instead of a race to the bottom' in terms of confidentiality and for ethical integration and usage, says Miller, who has been working on a risk assessment for BACP. 'One of the challenges I was tasked



'If something is bothering me, and it's not convenient to speak to a friend or my partner, I sometimes use AI for support. It can be surprisingly insightful' Katja Baghai-Ravary

with was to look at every single part of the profession through the lens of, "How might AI impact here?" It seemed prudent to assess risks first, because we want to make sure that we are deliberately moving towards safe applications instead of getting carried away with excitement around new AI abilities.'

Potential

Hanley thinks there is much for us to learn from the way clients are using AI therapy. 'AI is there at all hours of the day, not just for 50 minutes once a week, which is what you now see more corporately minded companies starting to consider. For instance, they're asking how can we use what we're learning from the way that people are actually using these less conventional services to challenge traditions in therapy? As a profession we are guided by many traditions, and we don't always question whether or not they're something that remains beneficial for the individuals working with us. It could be that we start thinking about how we can offer a package of support, rather than just a weekly 50 minutes, in which people can access quality controlled information and resources, and the in-person session is just one part of the therapy ecosystem. From one point of view, outcome studies demonstrate that we've not really got that much better at offering therapy over the years. This could be the point at which we enhance what we are offering.'

Trauma and EMDR therapist Katja Baghai-Ravary who uses AI in some client sessions (see 'Case study: creating a safe place', page 23) recommends practitioners explore AI through personal reflection as a way of getting to know its capabilities and limits. 'It's safer and more ethical to use your own issues rather than client presentations, even if you remove identifying information. If I want a

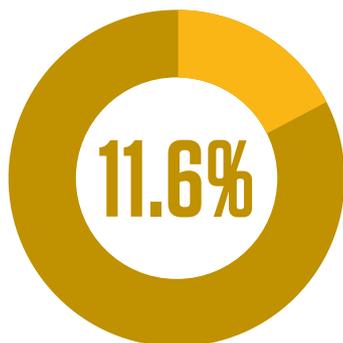
perspective on something that is bothering me, and it's not convenient to speak to a friend or my partner, I sometimes use AI for support. It can be surprisingly insightful. It's one reason why I believe that AI is going to influence our profession whether we're ready or not. I think rather than fearing it we need to approach it with awareness and intention.'

Siddals agrees that first-hand experience will help us best support clients in reflecting on their use of AI for support. 'Two participants in my study shared that their therapists seemed uncomfortable or confused on hearing that their clients were using AI between sessions, even though, ironically, one noted that his AI therapist was usually in agreement with his human therapist.'

Future developments

We have yet to see peak 'AI therapy' says Prescott, and she predicts that once we reach that we may see a return to prioritising therapy with a human. 'There is a lot of research on loneliness and how robots are increasingly used to support people. AI does build a bridge between human contact, helping people feel less lonely. But I think there will come a time when the realisation is that you're just talking to a robot and you actually really want human contact because we are social animals, aren't we? We want that social interaction.'

The next goal for the big tech companies is artificial general intelligence (AGI), with capability that matches human intelligence. If that happens, Miller says therapists may well be more in demand to help clients navigate the changes in our daily lives. 'If it's true that AI is coming for our jobs, then it's coming for everyone's jobs. Our role will be helping people who have suddenly gone from having "too busy" lives to not having enough to do.



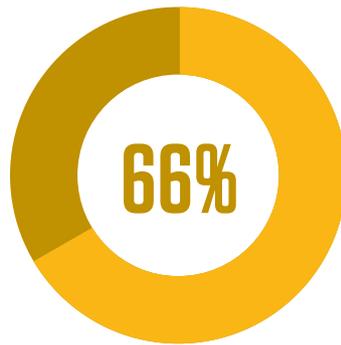
of BACP therapists are already using AI in their work⁶



We'll be working with people who need to think about who they are when their identity is no longer based on their profession. I think there's going to be plenty of work from whatever societal changes there might be.'

In the meantime, Miller says there are plenty of reasons to be positive about the potential uses for AI. 'We're the first generation who have AI and my feeling is that by raising the common bar of awareness about this, we can be confident in finding our own voice about the rights and wrongs of AI use. There is room for debate within our profession, and the more people who know about it I think the better the debates will start to be in the next few years. As practitioners we need to be informed enough about AI so that when clients mention to us they use it, we have a sense of whether their use is adaptive or maladaptive,' he says.

'The potential applications really are very exciting, and we should be hopeful about it but only once we know that we're safe and can demonstrate we're able to use it



of 25- to 34-year-olds would prefer to talk about their feelings with AI rather than a loved one⁸

well without harming ourselves or our clients. We need critical thinking and learning, and to train ourselves up on this so we keep asking questions. It's our challenge to set the standards today for how AI gets used in future.' ●

AI and the environment

One of the most compelling, but perhaps the least talked about, arguments for limiting our reliance on AI is the environmental impact, for several reasons:

- Both the development and the use of generative AI require significant amounts of electricity with resulting carbon emissions
- The microchips that power AI need rare earth elements, which are often mined in environmentally destructive ways
- All 'cloud computing' requires real-world infrastructure in the form of data centres housing servers, data storage drives and network equipment. Those that house generative AI require seven or eight times more energy to run than normal computer data centres
- Large amounts of water are needed as part of the cooling process during the development of AI, putting pressure on local water supplies
- Every usage of ChatGPT requires an estimated five times more energy than a normal internet search.¹⁶

- Richard Miller runs online 'AI and the human connection' workshops for BACP. To book onto the next one, see baccp.co.uk/events 🖱️

REFERENCES

1. Zao-Sanders M. How people are really using gen AI in 2025. Harvard Business Review 2025; 9 April. bit.ly/4dolwYF
2. Heinz MV et al. Randomized trial of a generative AI chatbot for mental health treatment. The New England Journal of Medicine AI 2025; 2(4).
3. Maurya RK et al. Assessing the use of ChatGPT as a psychoeducational tool for mental health practice. Counselling and Psychotherapy Research 2025; 25: e12759.
4. Yonatan-Leus R, Brukner H. Comparing perceived empathy and intervention strategies of an AI chatbot and human psychotherapists in online mental health support. Counselling and Psychotherapy Research 2025; 25: e12832.
5. Yin Y, Jia N, Waksak CJ. AI can help people feel heard, but an AI label diminishes this impact. Proceedings of the National Academy of Sciences of the United States of America 2024; 121(14).
6. BACP. Technology in therapy survey. Lutterworth: BACP; 2024.
7. Prescott J, Hanley T. Therapists' attitudes towards the use of AI in therapeutic practice: considering the therapeutic alliance. Mental Health and Social Inclusion 2023; 27(2).
8. The University of Law. Lonely millennials choose AI over a loved one to share their problems. [Press release.] 2023; 2 November. bit.ly/4jQiztE
9. Siddals S, Torous J, Coxon A. 'It happened to be the perfect thing': experiences of generative AI chatbots for mental health. npj Mental Health Research 2024; 3: 48.
10. He Y et al. Conversational agent interventions for mental health problems: systematic review and meta-analysis of randomized controlled trials. Journal of Medical Internet Research 2023; 25.
11. Ahmed A et al. Thematic analysis on user reviews for depression and anxiety chatbot apps: machine learning approach. JMIR Formative Research 2022; 6(3).
12. OpenAI. Sycophancy in GPT-4o: what happened and what we're doing about it. [Press release.] 2025; 29 April. openai.com/index/sycophancy-in-gpt-4o
13. Abrams Z. Using generative AI chatbots for mental health support: a dangerous trend. American Psychological Association 2025; 12 March.
14. Bridgeman J, Giraldez-Hayes A. Using artificial intelligence-enhanced video feedback for reflective practice in coach development: benefits and potential drawbacks. Coaching: An International Journal of Theory, Research and Practice 2023; 17(1): 32-49.
15. Prescott J, Ogilvie L, Hanley T. Student therapists' experiences of learning using a machine client: a proof-of-concept exploration of an emotionally responsive interactive client (ERIC). Counselling and Psychotherapy Research 2024; 24(2): 524-531.
16. Zewe A. Explained: generative AI's environmental impact. MIT News 2025; 17 January. bit.ly/43ikeuE

ABOUT THE AUTHOR

Sally Brown MBACP is a therapist and coach in private practice, and a freelance journalist and editor.

