

INSIGHTS / HUMAN-CENTRED AI

The AI Build Race Is Missing the Human Brain

Stephen Fahey

Founder

Steve's AI Learning & Development Services

September 2026

Executive Summary

The current AI market is moving at extraordinary speed. Founders are building products, organisations are experimenting with agentic workflows, and businesses are searching for competitive advantage wherever they can find it. Much of that attention is understandably focused on models, proprietary data, automation and technical capability. However, there is another side of the equation that I believe deserves far more attention: the human being expected to use all of this technology. AI can become more powerful, more autonomous and more sophisticated, but people still have limited cognitive capacity. We become tired, distracted and overloaded. We need time to learn, interpret, question and make decisions. For me, the future of AI adoption therefore cannot simply be about what technology can do. It must also consider whether humans are actually ready to orchestrate it safely and effectively.

We Are Building Faster Than Humans Are Learning

One of the things I keep noticing across the AI market is just how much emphasis is being placed on building. Companies are building products, startups are building agents, developers are building workflows, and organisations are experimenting with new ways of automating work. There is absolutely nothing wrong with that. Innovation requires experimentation. My concern is that we may be building considerably faster than we are thinking about how the human head actually works.

That distinction matters because an AI product does not operate in isolation. Eventually somebody has to use it. Somebody has to understand the interface, interpret its recommendations, decide when to trust it, recognise when something looks wrong and understand when human judgement needs to take over. As AI becomes increasingly agentic, that relationship may become even more complicated. Rather than simply using software, employees may increasingly find themselves orchestrating several AI systems that are completing different tasks, passing information between one another and making recommendations on their behalf.

My own working hypothesis is that this could place considerable pressure on what I describe as **brain capital**. Humans only have so much attention and cognitive energy available during a working day. Giving somebody five AI tools does not automatically make them five times more productive. In some environments it could potentially create more monitoring, more checking, more context switching and more decisions. That does not mean agentic AI will fail. It means we need much better evidence about how human beings actually respond when these systems become part of everyday work.

AI Fluency Is More Than Knowing How to Prompt

This is where I believe AI fluency needs a broader definition. Being fluent with AI should not simply mean knowing how to write a good prompt or operate the latest platform. It should include understanding when AI is useful, when it should be questioned, how its outputs should be evaluated and how to maintain your own critical thinking while using it.

Every industry will experience this differently. A lawyer working with an AI research system will face very different cognitive demands from somebody using AI in marketing, healthcare, finance, recruitment or manufacturing. Even two companies operating in exactly the same industry may adopt AI in very different ways. Their data, culture, workflows, leadership and employee capability will all influence the outcome. That is why I struggle with universal claims about AI adoption. There is unlikely to be one model of human readiness that works everywhere.

Companies understandably want to create leverage quickly. Proprietary data can give businesses significant advantages, and well-designed AI systems may dramatically improve certain workflows. However, competitive advantage cannot simply be built into the technology. It also needs to exist in the people using it. An organisation could purchase an extraordinary AI platform and still achieve relatively little if its employees do not understand how to use it effectively.

The Learning and Development Challenge Is Getting Bigger

This creates an important challenge for learning and development. The workforce entering organisations over the next few years will have grown up with much greater access to AI than previous generations. That could produce people with exceptional technological capability. It could also create different weaknesses if individuals become accustomed to delegating research, writing, analysis and problem-solving to machines before developing those abilities deeply themselves.

I am particularly interested in what this could mean for critical thinking, empathy, communication and independent judgement. I do not think we should simply assume these capabilities will disappear, but neither should organisations assume they will automatically remain strong. If AI increasingly performs parts of the intellectual process during education and early careers, companies may eventually need to invest much more deliberately in developing the human skills required to supervise and challenge AI systems.

That means learning design becomes strategically important. Businesses might raise Series B, C or D funding, hire excellent engineers and build impressive technology, but they will still have employees who need to learn. They will still need leadership development. They will still need people who understand how the company works, what it stands for and how decisions should be made. Funding does not remove the human requirement for learning.

The same principle applies internally to AI startups themselves. Founders can become so focused on the product that they forget their own organisation is also evolving. The company grows, new colleagues arrive, responsibilities become more complicated and suddenly the organisation needs systems for knowledge transfer, capability development and human coordination. The technology may scale quickly. People usually require something different.

We Need to Test AI Through Human Eyes

This is becoming an increasingly important part of the thinking behind **Steve's AI Learning & Development Services**. Rather than trying to compete with engineers or become another AI development company, I am increasingly interested in what happens after somebody builds the technology.

Can an intelligent but non-technical person understand it? How much cognitive effort does the workflow require? Where does the user need to stop and think? When could somebody become over-reliant on the AI? What training would be required before an employee used the system confidently? What happens to the existing job when the technology is introduced? What questions should somebody be asking about the data being used and where that data is going?

Those are the kinds of questions I naturally find myself asking when I test new AI products. My background in learning and development influences that considerably. Coming from the classroom, I tend to look at systems through the eyes of the learner. Does the sequence make sense? Has something been assumed that has never been explained? Is the person being given too much information at once? Does the design encourage understanding, or simply encourage the user to keep clicking?

That perspective could become increasingly valuable as AI products become more sophisticated. Technical capability tells us whether something can work. Human-centred evaluation helps us understand whether people can realistically work with it.

The Opportunity Is Human Readiness

This is why I am becoming interested in two ideas: a **Human Adoption Diagnostic** and a **Human Readiness Review**. The first would examine the product itself from a human perspective. The second would look more broadly at what people and organisations need in order to adopt the technology responsibly and effectively.

I see this almost like a pre-check before a more formal assessment. When you take a vehicle for an MOT, there are basic things you might check beforehand: lights, tyres, fluids and obvious faults. That does not replace the MOT, and neither should a human readiness review pretend to replace specialist cybersecurity, legal, regulatory or AI governance expertise. Its value could instead be identifying obvious human and organisational risks before an organisation reaches those later stages.

That distinction is important. I do not want our work to become another attempt to claim expertise in every area of AI. The value may come from doing something much narrower and doing it consistently: asking whether the human beings surrounding the technology are genuinely ready for it.

There is still a great deal I need to learn, test and validate. That is exactly why I want to work with founders who are creating serious AI products and experience those systems firsthand. The aim is not to criticise what they have built. It is to contribute another perspective.

Because while the AI industry continues to ask, **“What can we build next?”**, I think somebody also needs to keep asking, **“What will all of this require from the human brain?”**