

Safe AI doesn't guarantee safe use

Juwon Akinyande

We're spending a lot of time asking whether AI is ready for healthcare.

Is it safe? Is it accurate?
Is the data protected?
Has it been properly evaluated?

These are important questions. We should be asking them.

But there is another question I don't think we're asking enough.

Are we ready to use it well?

I'm an occupational therapist working in the NHS. My interest in AI started outside healthcare. I was using the tools, learning how other people were using them and seeing how quickly they were becoming part of everyday work.

It made me wonder what this would mean when the same technology reached healthcare.

So earlier this year, I started Clinically Intelligent, a clinician-led

initiative exploring better ways to learn and work in healthcare.

Since then, I've delivered practical AI training to clinical staff. What I've learnt is that people aren't particularly interested in AI for the sake of AI.

They want to know whether it can make their work easier and help them provide better care.

The questions people ask me are practical.

Can I put this information into it?

How do I know whether what it gives me is right?

What should I actually use it for?

When shouldn't I use it?

Does using it actually make this piece of work any better?

These aren't really questions about prompting.

They're questions about judgement.

And this is where I think we're getting ahead of ourselves.

We can procure an approved system, put governance around it, activate thousands of licences and train people how to use it.

A licence is not a plan. Training isn't adoption either.

Someone can know how to write an excellent prompt and still give a system information they shouldn't.

They can receive a polished answer and fail to notice that it's wrong.

They can save five minutes generating something and spend ten minutes checking it.

They can become very good at using AI for work that never needed AI in the first place.

Real capability is being able to look at the work in front of you and decide whether AI belongs there at all.

If it does, what can you safely give it?

How will you judge the output? What remains your responsibility?

And afterwards, did anything actually improve?

That's the approach I'm developing through Clinically Intelligent:

Start with the work. Develop capability. Measure what changes.

Because we should be more demanding about what we call successful AI adoption.

Not licences activated.
Not courses completed.
Not prompts submitted.

Did it save meaningful time? Did quality improve?

Did we create new problems?

Did the people using it retain the professional judgement the work depended on?

There is rightly a huge amount of work going into AI assurance, governance, procurement and patient safety.

Keep doing it.

But eventually, even the safest technology reaches a human being. Often a very busy one.

Safe AI does not guarantee safe use.

AI will keep getting easier to use.

Our ability to use it well needs to keep up.

The technology matters.

But the human part is the point.
