



Highlights from this issue

doi:10.1136/archdischild-2017-314684

Ian Wacogne

I've been following, with interest, the developments in artificial intelligence (AI) over the past few years. Depending on which author you read, we're either a few years, or a lifetime away from, or several lifetimes from being replaced in all we do. There are some really interesting examples of how specific AI—where the software learns a very precise task—is nearly as good as, or perhaps better than humans at some things. Driving a car is one attracting much attention, and the recent examples of programs beating humans at board games seem to be the start of this. However, general AI appears to be further off. We might, quite soon, trust software to drive us down a motorway, but that exact same software will likely be rubbish at making a good cup of tea.

While we're still in the phase where we need to use wetware—our brains—in meatspace—the physical world—it seems likely that we, as healthcare professionals, will need to accept a few things, which are mostly adaptations to the limitations of that wetware in meatspace. First, we'll need to try to learn and understand and use the best sort of information we can find. Enter evidence based medicine, guidelines, and so on. And second we'll need to try to cope with the fact that

we treat the really messy, extraordinarily complex other inhabitants of this space, alongside lots of other folk dealing with the same constraints. Meaning: Humans, and working with humans, is tricky, and confounding, and amazing.

In the last couple of decades huge efforts have been put into developing the best possible evidence for us. We're lucky to be able to feature articles which draw on these in this journal from time to time. Sometimes the evidence is formalised as a national guideline—in the UK the National Institute for Health and Care Excellence issues quite a lot of them. And yet we don't really use them—not as much as we ought to. Runnacles, Roueché and Lachman look about why this is the case—why, as they quote, we only follow best practice guidelines around half of the time (*see page 27*). It's easy to assume lots of negative reasons for this—that folk are dim, or naughty, or just like to behave egregiously. These negative thoughts persist about as long as it takes us to think of the last time we ourselves deviated from a guideline—when of course we recall that we acted from high principle and in the absolute best interests of that patient in front of us. The general reality is that failure to follow guidelines is messy and complex and interesting.

These authors give us some helpful thoughts and tips about how we might improve perseverance with guidelines, and for this reason this is my editor's choice this month.

I'm not sure where editing a journal fits on the difficulty spectrum, but it is also error prone. We make a few, and when I say we, I mean "I" because I'm supposed to spot them. This is why I'd like to apologise to Kate Harvey for missing her name off the Azithromycin Picket in the last edition. In partial mitigation, you'll understand that Pickets are odd—we often have a different person writing the abstract from the person writing the commentary; in this case Kate wrote the abstract. But, I/we should have got it right, so sorry Kate.

On that note—and while my various tasks in the world are still mine and not yet taken over by various self learning, massively bootstrapping AIs—the really exciting and terrifying idea that an AI cleverer than us will be able, maybe very quickly, to program itself to be even cleverer and even cleverer than that - if you'd like to, get involved in writing, or get someone to come to talk to you group about how to write, then please do get in touch.

ianwacogne@nhs.net