

## Kaleidoscope

Derek K. Tracy, Dan W. Joyce,  
Sukhwinder S. Shergill

**Gender and research;** we recognise differences between men and women, but female participants are commonly under-included as trial participants. Sugimoto *et al* undertook a bibliometric analysis of gender-related reporting in 11.5 million research papers published between 1980 and 2016.<sup>1</sup> Over time, such reporting grew from 59% to 67% in clinical medicine, and from 36% to 69% in public health, but remained somewhat static, at around 31%, in biomedical research. We were surprised to learn that one ‘key’ reason for this was presumed female participant variability because of menstrual cycles; something demonstrated to be a myth (indeed, males have greater variability on many traits). An algorithm found publications with female first or last authorship did considerably better in terms of reporting: a particularly male blind-spot.

Interestingly this work also showed that better gender-related reporting was associated with publication in lower impact journals, which segues onto work by Witteman *et al.*<sup>2</sup> Men have long been shown to get more research funding than women, especially women suffering a confluence of societal disadvantage through being from black and minority ethnic backgrounds. However, the meaning of this observational finding has been controversial – for example, is it confounded by differences in the research proposals? The Canadian Institutes of Health Research tackled this by dividing funding applications into two streams, one of which explicitly assayed the calibre of the principal investigator. Overall grant application success was just over 15%, with women doing worse than men after adjusting for age and research domain. Crucially, when the two streams were compared, the difference was because of less favourable assessment of the women principal investigators, not the quality of their scientific proposal. It is 2019 folks, and just not good enough – #WomenInSTEM.

**A landmark 2016 paper<sup>3</sup> calculated that medical error was the third leading cause of death in the USA, accounting for an astonishing quarter of a million hospital deaths there annually.** This had an enormous impact, achieving coverage by a host of media outlets. Shojania & Dixon-Woods challenge the headline figure.<sup>4</sup> They note how it was derived from an average of several estimates from varying sources, further extrapolated to sample populations not covered in the original studies. Furthermore, the data were not rigorously challenged, not least the ability to unpick confounders and determine any causal role for common adverse events. They cite the example of an individual developing an allergic reaction to an antibiotic while in an intensive care unit with progressive multiorgan failure: this is unlikely to be the cause of death, but such nuance is difficult or impossible to unpick in most work. A key point is that ‘many patients die with, rather than of, these conditions’. This paper argues that the true figure of hospital deaths as a result of medical error is an order of magnitude lower: 3.6%, not 36%. Still too high, and none of this is to downplay the importance of patient safety; rather, given the understandable attention it attracts, it is to better know the real data.

Writing in the *New England Journal of Medicine*, Dinah Miller gives a thoughtful personal account<sup>5</sup> of an issue that has hit very many of us, but about which little is often said publicly: the sorrow we feel when one of our patients dies by suicide. She notes how data suggest half of psychiatrists will lose at least one patient this way during their career; while we rightfully think about the tragedy of the individual who has died, and the complex grief it can leave their loved ones, we tend not to discuss how it has an impact on us.

Doctors’ training prepares us for death, but there is something unique about death through suicide that hits us in a way death from physical health causes seldom does. It can never be an expected outcome, it is inevitably coloured with a sense of failure, and it invariably leads to questions of how different decisions might have altered outcomes. We suspect few ask that last question as harshly as psychiatrists do of themselves, and we recognise the tension of how to return to ‘normal practice’ through the self-recrimination. Dr Miller notes how we have no systematised way to come together, no ‘rituals of our own to mark a death and find a path toward healing’.

**We have recently reported on some disappointing trial results of interventions for suicide risk, but King *et al*<sup>6</sup> offer some hope with an intervention in adolescents.** They investigated the impact of a ‘youth-nominated support team’ – a psychosocial support of (non-mental health professional) ‘caring adults’ nominated by adolescents aged 11 to 14 who had been admitted to hospital following self-harm or significant suicidal thinking. On average this team encompassed about three such individuals, who came from varying environments including family, school and the local community. These supporters were given a training session to discuss the young person’s problem list and suicide warning signs, treatment plan and psychoeducation on communicating with adolescents and supporting positive behavioural choices. They also received weekly phone calls of support from the research team for a 3-month period. Compared with those randomised to treatment as usual, the group who received this active intervention had a 6.6-fold decrease in mortality in the 11 to 14 years after the initial hospital admission episode. Evidence of effectiveness over an enduring time period; although the nature of the study meant it was not possible to explore several important but non-fatal outcomes.

**Following the trumpeted National Health Service ‘Topol Review’,<sup>7</sup> the same author has assayed artificial intelligence in medicine,<sup>8</sup> claiming ‘Almost every type of clinician, ranging from specialty doctor to paramedic, will be using AI technology, and in particular deep learning, in the future’.** It is becoming increasingly hard to find a coherent definition of either artificial intelligence or machine learning that demarcates the methodological boundaries of either, or separates them from more familiar statistical methods. Tom Mitchell’s oft-cited definition of machine learning is ‘A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E’. Here, you can substitute T for predicting the probability of a disorder, outcome or prognosis; P measures the number of errors (cases predicted incorrectly) or how these match a ground-truth (clinician assessment). E is where the action is – experience implies iteration and error correction. At the most fundamental level, the program for finding the optimal (least squares) fit for a linear regression on some data is an iterative algorithm, minimising errors (the sum of squares of residuals) by performing operations on matrices representing cases and variates.

Topol interprets artificial intelligence broadly, but also somewhat specifically, by emphasising successes using deep neural networks. These implement an inherently hierarchical data-processing pipeline ideally positioned to ‘compress’ huge data (such as images) into a parsimonious set of ‘features’ that have the highest utility for making predictions (‘this chest x-ray shows a cancerous lesion in the left apical area’). Topol highlights that the strongest demonstrations have been in exactly these medical domains; where image processing (that would require human experts to use visual inspection) forms the basis of the diagnostic task. There is noticeable emphasis on super-human performance, but what is lacking is demonstration of improved outcomes.

For psychiatry, Topol cites examples of predicting suicide that had impressive headline performance. One study used electronic health records for predicting non-fatal suicide events, with the best performance being at 7-day follow-up, using an algorithm called random forests. In the clinical context, the reported precision and recall scores were 0.79 and 0.95 over 3250 positive (non-fatal attempts) and 1917 negative (controls) cases. What does this mean for clinical practice? This implies a false-negative rate of 162 cases of missed non-fatal suicide attempts and 820 false positives – that is, people predicted by the algorithm to make a non-fatal suicide attempt, but who did not. What did artificial intelligence add to practice?

**Neuroscience is data-rich but theory-poor – we have multitudinous experiments across species, but few theories that unify the diversity of findings.** One of the most challenging problems is how neurons code communications. Warren & McCulloch's landmark work hypothesised that they communicate using principles similar to a digital logic gate; specifically, neurons 'fire' or 'spike' with binary 0/1 outputs in response to their inputs. We know (by design) how digital computers achieve coherent communication – logic gates take 0/1 inputs (bits, or binary digits), compute something (operations such as AND, OR, XOR) and then 'output' results as further bits, which are passed to other logic gates connected to them. Most importantly, digital computers do this in a tightly synchronised fashion – all these logic gates compute and shunt bits around on the 'ticks' and 'tocks' from a central clock signal. It transpired that neurons behave differently. Their conveying of putative 'data' is not simply the presence (a one) or absence (a zero) of a 'spike' output, but rather, variation in their spike rates/frequencies, magnitudes and the relative timing/phases of trains of spikes all appear to be important in conveying or coding different features of stimuli or controlling the activity of the organism.

Pryluk *et al* propose a principle for neural coding as a robustness–efficiency trade-off.<sup>9</sup> Given a physiological neuron has an upper limit on the number of spikes it can produce, efficiency is the amount of information contained in the observed spike train compared with a theoretical maximum. Robustness is the correlation of pairs of individual or populations of neurons – so that patterns of spike trains that are strongly correlated are deemed to be robustly responding to a given stimulus. To compute efficiency, they use the notion of entropy – Claude Shannon's measure of the 'surprise' or unexpectedness in a signal transmitted over a channel with finite bandwidth. To use entropy in this way, one has to divide continuous quantities (time, number of spikes) into discrete 'bins'. Using an orthographic analogy, they define a discrete 'letter' by dividing spike trains into spikes occurring in time windows (i.e. a spike at 1, 2, 4, 8 and 16 ms) and a 'word' as the number of letters (4, 8 and 16 letters). Then, they are able to define the entropy of the 15 letter–word combinations by (a) measuring physiological neurons' spike trains and computing the probability of letter–word combinations and (b) simulating neurons (with the same firing rate as the physiological neurons) to arrive at theoretical upper and lower limits for the entropy. The divergence between the theoretical entropy limit and the observed (physiological) entropy is the basis of their analyses.

The experiments they report use cellular recordings from the amygdala and cingulate cortex in five macaque monkeys and seven individuals with pharmacologically intractable epilepsy. Using *in vivo* recordings for isolated, pairs and triplets of neurons, they compared the efficiency–robustness properties across primate species and cortical and subcortical structures. Their results showed that, in monkeys and humans, the neural code was more efficient (exploiting the theoretical communication limits implied by entropy) in the cortical structure than the amygdala, which they propose subserves

flexible cognition. In the amygdala, however, they found more robust (rather than efficient) neural coding, which, tentatively, supports the evolutionary role of the amygdala for threat detection (for example where a robust – rather than flexible – but not always accurate or efficient response is required). Further, the respective efficiency and robustness measures were higher in the human than primate for both cingulate and amygdala, which corresponds to the evolutionary lineage and cognitive capacities of the two species. They suggest that the evolutionary path (from lower-primate to human) and the preserved efficiency–robustness trade-off leads to a highly flexible human brain with the unfortunate consequences that it can over-adapt to amygdala and limbic stimuli, resulting in anxiety disorders including post-traumatic stress disorder.

**Finally, 'we don't even know what 90% of the brain does!' is oft-repeated nonsense, but following on from the last piece, an intriguing review paper argues<sup>10</sup> that the majority of neurons never fire and are 'permanently silent'.** Many electrophysiological data-sets show that most neurons do not produce action potentials when stimulated – up to 90% in some animal studies. The burden and inefficiency of having so many cells maintained in a quiescent state without obvious gain seems peculiar, but Ovsepián argues that these dormant cells, which he labels the 'dark matter of the brain' are phylogenetically ancient evolutionary remnants whose inactivity makes them unsusceptible to the pressures of natural selection. However, he argues that stress and illness can activate them, and this is a driver for some neuropsychiatric conditions, including Tourette syndrome, autism spectrum disorders and psychoses. Big calls need big evidence, and we found the argument fascinating but ultimately unconvincing. The respected 'neuroskeptic' blog in *Discover* provides an interesting counter<sup>11</sup> to the paper, addressing, among other issues, the concept of 'sparse firing' where most neurons will not react in response to a stimulus; however, this argued principle of brain organisation does not mean they cannot or will not react to the appropriate input. Fascinating stuff, and we leave the last word to Ovsepián, who reminds us of Freud's line 'where does a thought go when it is forgotten?'

## References

- 1 Sugimoto CR, Ahn Y-Y, Smith E, Macaluso B, Larivière V. Factors affecting sex-related reporting in medical research: a cross-disciplinary bibliometric analysis. *Lancet* 2019; **393**: 550–9.
- 2 Witterman HO, Hendricks M, Straus S, Tannenbaum C. Are gender gaps due to evaluations of the applicant or the science? A natural experiment at a national funding agency. *Lancet* 2019; **393**: 531–40.
- 3 Makary MA, Daniel M. Medical error—the third leading cause of death in the US. *BMJ* 2016; **353**: i2139.
- 4 Shojania KG, Dixon-Woods M. Estimating deaths due to medical error: the ongoing controversy and why it matters. *BMJ Qual Saf* 2017; **26**: 423–8.
- 5 Miller D. When a patient dies by suicide – the physician's silent sorrow. *NEJM* 2019; **380**: 3111–4.
- 6 King CA, Arango A, Kramer A, Busby D, Czyz E, Foster CE, et al. Association of the youth-nominated support team intervention for suicidal adolescents with 11- to 14-year mortality outcomes secondary analysis of a randomized clinical trial. *JAMA Psychiatry* 6 Feb 2019 (doi:10.1001/jamapsychiatry.2018.4358).
- 7 Health Education England. *The Topol Review*. HEE, 2019 (<https://www.hee.nhs.uk/our-work/topol-review>).
- 8 Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nature Med* 2019; **25**: 44–56.
- 9 Pryluk R, Kfir Y, Gelbard-Sagiv H, Fried I, Paz R. A tradeoff in the neural code across regions and species. *Cell* 2019; **176**: 597–609.
- 10 Ovsepián SV. The dark matter of the brain. *Brain Struct Funct* 18 Jan 2019 (doi: 10.1007/s00429-019-01835-7).
- 11 Neuroskeptic. Silent neurons: the dark matter of the brain? *Discover* 6 Feb 2019 ([http://blogs.discovermagazine.com/neuroskeptic/2019/02/06/silent-neurons-dark-matter-brain/#.XGWA\\_CHAOUm](http://blogs.discovermagazine.com/neuroskeptic/2019/02/06/silent-neurons-dark-matter-brain/#.XGWA_CHAOUm)).