Don't govern AI with policies designed for drugs

Health authorities are overlooking risks to systems and society in their evaluations of new digital technologies, says Melanie Smallman

In many ways the code of conduct for artificial intelligence systems in healthcare just rolled out by the UK Government is timely, necessary, and likely to serve as a global benchmark. The principles, laid out by the Department of Health and Social Care, aim to protect patient data and "ensure that only the best and safest data-driven technologies are used." The sorts of projects they relate to include Deepmind crunching through more than a million eye scans with London's Moorfelds Eye Hospital to design an algorithm to detect macular degeneration and Ultromics and John Radcliffe Hospital in Oxford working with AI to improve detection of heart disease and lung cancer.

Yet I fear the new guidelines might be the start of a deluge of inadequate policies to regulate AI. The guidelines neglect how AI is changing the health care system, community and society [Pull quote], instead adhering to traditional assessments on medical interventions and impacts on individual privacy, safety and efficacy.

The impact of AI is more akin to that of automobiles or personal computers than to medicine. Medicines are prescribed to patients; their use closely tied to individual need. Cars have shaped all our lives, our cities, and industries, even for individuals who do not drive. Policy around innovation and technology largely ignores tech's potential to worsen inequalities, even as examples mount. US sociologist Virginia Eubanks recently coined the phrase 'digital poorhouse' for the effects of AI and automation on low-income households and communities. For example, in the City of Los Angeles, which uses a program to match homeless people with the most appropriate housing available, to gain shelter, individuals are asked to state their names, whether they have had unprotected sex or considered self-harm, and how often they accessed crisis services for sexual assault or domestic abuse. Middle class communities would not tolerate this level of intrusion. And when data like this is coordinated between social services, police and other services, the potential for new forms of unfairness is wide open.

I study the relationship between science and society at University College London and am part of a team considering Data Ethics and AI in health care at London's Alan Turing Institute. The power to widen inequality is not merely a consequence of slipshod regulation or unintended side effects: it is deeply embedded in technologies themselves. For instance, the success of most digital businesses depends upon producing and selling goods without huge manufacturing and distribution costs. This raises salaries for high-skill workers while reducing demand and conditions for lower skill jobs.

We can already glean how new technologies are changing healthcare systems. Last year London-based Babylon Health's new app, which provided smartphone GP consultations, was criticised by the Royal College of General Practitioners for cherry picking patients, leaving traditional GP services to deal with the most complex patients, without sufficient resources to do so; Anthony Zietman from Massachusetts General Hospital has described how the huge expense of proton beam therapy units distort healthcare markets in the US and channel funds from other areas of need, such as traditional radiotherapy. My colleagues at Kings College have found that investment in surgical robotics draws funds from other

treatments and centralises care in large teaching hospitals, requiring many patients to travel longer distances or forego treatment.

The public understands that the pros and cons of technologies are often inextricably linked, that evaluating technologies means deciding whether benefits outweigh the downsides, and that doing so depends on how both are distributed. Over more than a decade of using focus groups and participatory exercises to gauge public opinion—on topics from stem cells to nanoscience -- I have seen consistently sophisticated public assessments of how effects are felt at multiple, interacting scales, from individual to society. People worry about the kind of world technologies will create, not just the risk of harm to individuals. Our policies must show similar levels of sophistication.

In my view, the new code is a missed opportunity to get things right the first time, to anticipate the wider problems that are inevitable, and to keep the health system affordable and effective. It is thanks to the affordable, effective, equitable and comprehensive health service that the UK has more than seven decades of data - crucial for developing AI for healthcare. But this same data has also helped show that social inequality is detrimental to the physical and mental health of all members of society -- through increased stress, with documented biological effects ranging from chronic inflammation, to chromosomal aging and brain function.

Health technologies can improve health care, reduce costs, increase the speed of diagnostics, and benefit society. But fulfilling that potential will require us to broaden the lens through which we evaluate them, and soon.

Doing so will not be simple. As with the advent of the motor-car, many of their most serious implications will be emergent and the harshest effects likely to be felt by communities with the least powerful voices. We need a framework that forces us to move our gaze from individuals to systems to communities, and back again. This means bringing together different types of expertise – including that of workers and citizens - to develop a framework that the health systems can use to anticipate and address these issues on a case by case basis. Specifically, this framework needs to have an explicit mandate to consider and anticipate the social consequences of AI in healthcare – and to keep a watching eye on these effects in the future. That is the best way to ensure new health technologies meet the needs of everybody, and not just those in silicon valley.

Melanie Smallman is Co-Director of the Responsible Research and Innovation Hub at University College London's Department of Science and Technology Studies and a Fellow at the Alan Turing Institute. m.smallman@ucl.ac.uk