theguardian.com

What does it mean to be human in the age of technology?

Tom Chatfield

9-11 minutes

When I think about the future of human-machine interactions, two entwined anxieties come to mind.

First, there is the tension between individual and collective existence. Technology connects us to each other as never before, and in doing so makes explicit the degree to which we are defined and anticipated by others: the ways in which our ideas and identities do not simply belong to us, but are part of a larger human ebb and flow.

This has always been true – but rarely has it been more evident or more constantly experienced. For the first time in human history, the majority of the world's population is not only literate – itself an achievement less than a century old – but also able to actively participate in written and recorded culture, courtesy of the connected devices pervading almost every country on earth.

This is an astonishing, disconcerting, delightful thing: the crowd in the cloud becoming a stream of shared consciousness.

Second, there is the question of how we see ourselves. Human nature is a baggy, capacious concept, and one that technology has

altered and extended throughout history. Digital technologies challenge us once again to ask what place we occupy in the universe: what it means to be creatures of language, self-awareness and rationality.

Our machines aren't minds yet, but they are taking on more and more of the attributes we used to think of as uniquely human: reason, action, reaction, language, logic, adaptation, learning. Rightly, fearfully, falteringly, we are beginning to ask what transforming consequences this latest extension and usurpation will bring.

I call these anxieties entwined because, for me, they come accompanied by a shared error: the overestimation of our rationality and our autonomy. In asking what it means to be human, we are prone to think of ourselves as individual, rational minds, and to describe our relationships with and through technology on this basis: as isolated "users" whose agency and freedom are a matter of skills and reasoned options; as task-performers who are existentially threatened by any more efficient agent.

This is one view of human-machine interactions. Yet it's also an account of human beings that gives us at once too little and too much credit. We know ourselves to be intensely social, emotional, intractably embodied creatures. Much of the best recent work in economics, psychology and neuroscience has emphasized the degree to which we cannot be unbundled into distinct capabilities: into machine-like boxes of distinct memory, processing and output.

Neither language, culture nor a human mind can exist in isolation, or spring into existence fully formed. We are interdependent to an extent we rarely admit. We have little in common with our creations

– and a nasty habit of blaming them for things we are doing to ourselves.

What makes all this so urgent is the brutally Darwinian nature of technological evolution. Our machines may not be alive, but the evolutionary pressures surrounding them are every bit as intense as in nature, and with few of its constraints. Vast quantities of money are at stake, with corporations and governments vying to build faster, more efficient and more effective systems; to keep consumer upgrade cycles ticking over. To be left behind – to refuse to automate or adopt – is to be out-competed.

As the philosopher Daniel Dennett, among others, has pointed out, this logic of upgrade and adoption extends far beyond obvious fields such as finance, warfare and manufacturing. If a medical algorithm is proven to produce more consistently accurate diagnoses than a physician, it's both unethical and legally questionable to refuse to use it. As self-driving or semi-autonomous cars become more affordable and road-legal, it's hard to argue against the ethical and regulatory case for making them obligatory. And so on. Few fields of human endeavour are likely to remain untouched.

Machines, in other words, are becoming stunningly adept at making decisions for us on the basis of vast amounts of data – and getting better at this at an equally stunning rate. Forget the hypothetical emergence of general purpose artificial intelligence, at least for a moment: we're handing over more and more of what happens in our world, today, to the speed and efficiency of unthinking deciders.

It's precisely because our present machines can neither think nor feel that this matters. We call them "smart" and marvel at their

powers; we paint pictures of a world in which they, not we, are determining what we do and how. We can't help ourselves: we see purpose, autonomy and intent everywhere.

Yet in ascribing agency and intentions to our tools that they don't possess, we misunderstand several fundamental points. Humans aren't slow, dumb and heading for the evolutionary scrapheap; machine efficiency is a very poor model indeed for understanding ourselves; and cutting people out of every possible loop – the better to assure speed, profit, protection or military success – is a poor model for a future in which humans and machines equally maximise their capabilities.

Our creations are effective in part because they are unburdened by most of what makes humans human: the broiling biological pot of emotion, sensation, bias and belief that constitutes the bulk of mental life. We are biased, beautiful creatures. Technology and intellect allow us to externalise our goals; but the ends pursued are those we chose.

Do the incentives our tools tirelessly pursue on our behalf include human thriving, meaningful work, rich and humane interactions? Do we believe these things to be unachievable, unknowable or worthless? If not, when are we going to shift our focus?

If we wish to build not only better machines, but better relationships with and through machines, we need to start talking far more richly about the qualities of these relationships; how precisely our thoughts and feelings and biases operate; and what it means to aim beyond efficiency at lives worth living.

What does a successful collaboration between humans and machines look like? One, I would argue, in which humans remain in

the loop, able transparently to assess a system's incentives – and either to influence its direction or debate its alteration.

What does a successful collaboration between humans mediated by technology look like? We have plenty of these already, and they're characterised by the maximisation of all resources involved: human creativity and questioning; machine search, speed, processing and recall; an iteration involving all parties; and the recognition that efficiency is not an end in itself, but simply a measure of velocity.

Finally, let's be clear about one thing. Ours is an amazing time to be alive: to be debating such questions together. If there's one thing our swelling collective articulacy as a species brings home, it's that people care above all about other people: what they think, do, believe, fear, hate, love, laugh at – and what we can make together.

Our creations are certain to grow far beyond our current comprehension: how far and how fast is perhaps our most urgent existential question. Our best hopes of progress, however, remain deceptively familiar: understanding ourselves better; asking what aims may serve not only our survival, but also our thriving; and striving to build systems that serve rather than subvert these.

- ***This is an edited extract from Tom Chatfields's address at the launch of the Humanities and Digital Age programme, led by*** [*The Oxford Research Centre in the Humanities*](#) ***at Oxford University in the UK, on 21 January. The discussion*** [*will be broadcast live*](#) ***on 21 January from 17.30 GMT and*** [*some free tickets are available*](#)

Since you're here…