mediastudies.press • Social Media & the Self: An Open Reader

How Much Can We Afford to Forget, if We Train Machines to Remember?

Gene Tracy

Published on: Apr 08, 2019

DOI: 10.32376/3f8575cb.c0097cd0

License: [Creative Commons Attribution-NoDerivatives 4.0 International License \(CC-BY-ND 4.0\)](https://creativecommons.org/licenses/by-nd/4.0/)

WHEN I WAS a student, in the distant past when most computers were still huge mainframes, I had a friend whose PhD advisor insisted that he carry out a long and difficult atomic theory calculation by hand. This led to page after page of pencil scratches, full of mistakes, so my friend finally gave in to his frustration. He snuck into the computer lab one night and wrote a short code to perform the calculation. Then he laboriously copied the output by hand, and gave it to his professor.

Perfect, his advisor said—this shows you are a real physicist. The professor was never any the wiser about what had happened. While I've lost touch with my friend, I know many others who've gone on to forge successful careers in science without mastering the pencil-and-paper heroics of past generations.

It's common to frame discussions of societal transitions by focusing on the new skills that become essential. But instead of looking at what we're learning, perhaps we should consider the obverse: what becomes safe to forget? In 2018, *Science* magazine asked dozens of young scientists what schools should be teaching the next generation. Many [said](#) that we should reduce the time spent on memorising facts, and give more space for more creative pursuits. As the internet grows ever more powerful and comprehensive, why bother to remember and retain information? If students can access the world's knowledge on a smartphone, why should they be required to carry so much of it around in their heads?

Civilisations evolve through strategic forgetting of what were once considered vital life skills. After the agrarian revolution of the Neolithic era, a farm worker could afford to let go of much woodland lore, skills for animal tracking, and other knowledge vital for hunting and gathering. In subsequent millennia, when societies industrialised, reading and writing became vital, while the knowledge of ploughing and harvesting could fall by the wayside.

Many of us now rapidly get lost without our smartphone GPS. So what's next? With driverless cars, will we forget how to drive ourselves? Surrounded by voice-recognition AIs that can parse the most subtle utterances, will we forget how to spell? And does it matter?

Most of us no longer know how to grow the food we eat or build the homes we live in, after all. We don't understand animal husbandry, or how to spin wool, or perhaps even how to change the spark plugs in a car. Most of us don't need to know these things

because we are members of what social psychologists [call](#) ‘transactive memory networks’.

We are constantly engaged in ‘memory transactions’ with a community of ‘memory partners’, through activities such as conversation, reading and writing. As members of these networks, most people no longer need to remember most things. This is not because that knowledge has been entirely forgotten or lost, but because someone or something else retains it. We just need to know whom to talk to, or where to go to look it up. The inherited talent for such cooperative behaviour is a gift from evolution, and it expands our effective memory capacity enormously.

What’s new, however, is that many of our memory partners are now smart machines. But an AI—such as Google search—is a memory partner like no other. It’s more [like](#) a memory ‘super-partner’, immediately responsive, always available. And it gives us access to a large fraction of the entire store of human knowledge.

Researchers have identified several pitfalls in the current situation. For one, our ancestors evolved within groups of other humans, a kind of peer-to-peer memory network. Yet information from other people is invariably coloured by various forms of bias and motivated reasoning. They dissemble and rationalise. They can be mistaken. We have learned to be alive to these flaws in others, and in ourselves. But the presentation of AI algorithms inclines many people to believe that these algorithms are necessarily correct and ‘objective’. Put simply, this is magical thinking.

The most advanced smart technologies today are trained through a repeated testing and scoring process, where human beings still ultimately sense-check and decide on the correct answers. Because machines must be trained on finite data-sets, with humans refereeing from the sidelines, algorithms have a tendency to amplify our pre-existing biases—about race, gender and more. An internal recruitment tool used by Amazon until 2017 presents a classic case: trained on the decisions of its internal HR department, the company found that the algorithm was systematically sidelining female candidates. If we’re not vigilant, our AI super-partners can become super-bigots.

A second quandary relates to the ease of accessing information. In the realm of the nondigital, the effort required to seek out knowledge from other people, or go to the library, makes it clear to us what knowledge lies in other brains or books, and what lies in our own head. But researchers [havefound](#) that the sheer agility of the internet’s

response can lead to the mistaken belief, encoded in later memories, that the knowledge we sought was part of what we knew all along.

Perhaps these results show that we have an instinct for the ‘extended mind’, an idea first [proposed](#) in 1998 by the philosophers David Chalmers and Andy Clark. They suggest that we should think of our mind as not only contained within the physical brain, but also extending outward to include memory and reasoning aids: the likes of notepads, pencils, computers, tablets and the cloud.

Given our increasingly seamless access to external knowledge, perhaps we are developing an ever-more extended ‘I’—a latent persona whose inflated self-image involves a blurring of where knowledge resides in our memory network. If so, what happens when brain-computer interfaces and even brain-to-brain interfaces become common, perhaps via neural implants? These [technologies](#) are currently under development for use by locked-in patients, stroke victims or those with advanced ALS, or motor neurone disease. But they are likely to become far more common when the technology is perfected—performance enhancers in a competitive world.

A new kind of civilisation seems to be emerging, one rich in [machine intelligence](#), with ubiquitous access points for us to join in nimble artificial memory networks. Even with implants, most of the knowledge we’d access would not reside in our ‘upgraded’ cyborg brains, but remotely—in banks of servers. In an eye-blink, from launch to response, each Google search [now](#) travels on average about 1,500 miles to a data centre and back, and uses about 1,000 computers along the way. But dependency on a network also means taking on new vulnerabilities. The collapse of any of the webs of relations that our wellbeing depends upon, such as food or energy, would be a calamity. Without food we starve, without energy we huddle in the cold. And it is through widespread loss of memory that civilisations are at risk of falling into a looming dark age.

But, even if a machine can be said to think, humans and machines will think differently. We have countervailing strengths, even if machines are often no more objective than we are. By working together in human-AI teams, we can play superior chess and make better medical decisions. So why shouldn’t smart technologies be used to enhance student learning?

Technology can potentially improve education, dramatically widen access, and promote greater human creativity and wellbeing. Many people rightly sense that they

stand in some liminal cultural space, on the threshold of great change. Perhaps educators will eventually learn to become better teachers in alliance with AI partners. But in an educational setting, unlike collaborative chess or medical diagnostics, the student is not yet a content expert. The AI as know-it-all memory partner can easily become a crutch, while producing students who think they can walk on their own.

As the experience of my physicist friend suggests, memory can adapt and evolve. Some of that evolution invariably involves forgetting old ways, in order to free up time and space for new skills. Provided that older forms of knowledge are retained somewhere in our network, and can be found when we need them, perhaps they're not really forgotten. Still, as time goes on, one generation gradually but unquestionably becomes a stranger to the next.

*[Gene Tracy](#) is chancellor professor of physics at William & Mary, Virginia. He is the author of *Ray Tracing and Beyond: Phase Space Methods in Plasma Wave Theory* (2014). He blogs about science and culture at [The Icarus Question](#).*

REPRINT

[“How Much Can We Afford to Forget, if We Train Machines to Remember?”](#) (Gene Tracy, *Aeon*, April 8, 2019)

OPENLY LICENSED