

ARTICLE · CULTURA & SAÚDE

The Productivity That Leaves No Memory

Memory, relationships and speed in the age of artificial intelligence agents

Dr. Alexandre Campos Moraes Amato

Vascular surgeon

Received for publication in September 2026 · Section: Article · Language: English

ABSTRACT

Drawing on his own experience with artificial intelligence agents at scale, the author examines three linked concerns that hyperproductivity brings to the surface. The first is memory: projects completed with the machine's help tend to be forgotten entirely, which is consistent with the generation effect, desirable difficulties and preliminary evidence of lower retention among users of language models. The second is human interaction: by taking over the practical function of everyday exchanges, technology removes the pretext for the encounter, a pattern suggested by data on loneliness and on the behavior of young people in the smartphone era, without establishing causality. The third is speed: technology advances faster than social adaptation, a situation described by the Collingridge dilemma and illustrated by the late evidence on screens and child development. He concludes that the response within the individual's reach is modest and domestic: deliberate caution, example, named rules and effort preserved where it still builds memory.

KEYWORDS: *artificial intelligence; memory; productivity; human relationships; child development; technology*

The list that never empties

Anyone who has started working with artificial intelligence agents at scale will recognize a recent phenomenon. Dozens of processes run in parallel, each on a different task, and individual productivity reaches levels that seemed impossible only a few years ago. Projects that would have slept in a drawer for years come out in days. Work that would have occupied a small team for months runs overnight, unsupervised, on a single computer.

The first surprise is not the speed but what it reveals. Ideas that seemed capable of changing everything turn out, once finally executed, to be niche: interesting to the person who conceived them and to almost no one else. The machine makes it possible to travel in hours down paths that would once have taken years, and to discover quickly that many of them lead nowhere. That has value, but not the value one imagined.

The second surprise is more uncomfortable. The to-do list, emptied with unprecedented efficiency, fills up again at the same pace. Each completed task suggests three more, and the tool itself, at the end of each session, proposes the next step. It feels like a bottomless pit. The maxim attributed to Paracelsus applies here: the difference between remedy and poison lies in the dose.

What follows is not a condemnation of artificial intelligence, which few of those who have tried it seriously would give up. It is an attempt to name three linked concerns that hyperproductivity brings to the surface, about memory, about human relationships and about the speed of change, and to point to the only response that seems within the individual's reach: a modest one, and one that fits within the home.

The memory of effort

The first concern arises from a comparison any researcher can make with themselves. Details of a thesis

defended fifteen years ago remain vivid, because they cost day after day of immersion. Of projects completed with the machine's help, sometimes larger than the thesis, one forgets not the detail but the entire project. The explanation is plausible from a neurobiological standpoint: we retain through repetition and effort, and whatever passes through us too quickly leaves no groove.

Cognitive psychology has known this principle for decades. The generation effect, described by Slamecka and Graf in 1978, shows that information produced by the subjects themselves is retained better than information merely read. Robert Bjork coined the term desirable difficulties for the conditions that make learning slower and more effortful in the moment, and for that very reason more durable. Generative artificial intelligence is a machine for removing difficulties. When it removes the desirable ones along with the useless ones, the price is paid in the memory of whoever uses it.

There is preliminary evidence that this is already happening. In 2025, a group at the MIT Media Lab followed 54 volunteers writing essays under three conditions: with a language model, with a search engine, or with their own minds alone. Those who used the model showed the weakest neural connectivity on electroencephalography, the lowest sense of authorship and, revealingly, difficulty quoting passages from their own text minutes after handing it in. The study is small, was released as a preprint and needs replication, but it describes with experimental precision what many users report through introspection.

This reconfiguration did not begin with artificial intelligence. Earlier generations knew dozens of telephone numbers by heart; today many people do not know the number of the person they share a home with, though they can still recite the phone numbers of their childhood. Anyone who works with neuropsychological assessment knows that test norms must be recalibrated periodically, because what counts as normal memory shifts along with habits. The so-called Google effect, described by Sparrow, Liu and Wegner in 2011, proposed that we stop retaining what we expect to be able to look up later, remembering where the information is rather than the information itself. In fairness, the original experiment was not replicated in a large 2018 effort, and later studies find the effect only under specific conditions. Human memory is not a container that empties; it migrates to other functions.

The question that remains is whether this loss matters. Forgetting telephone numbers has cost no one anything, except on the day the battery died. But telephone numbers are data; a completed project is experience. It is possible to imagine looking back twenty years from now and finding, in place of this period, a void: a stretch of intense production of which no memory remains because nothing cost enough to be remembered. It is not dementia in the

clinical sense. It is something new, for which there is not yet a name.

The interaction that evaporates

The second concern is social. Start with the solitary work par excellence, writing. In January 2024, the winner of the Akutagawa Prize, the most prestigious award in Japanese literature, revealed that about 5% of the prize-winning novel had been generated verbatim by a language model. The book is still hers, but the work of writing has become something else. Writing demanded prolonged engagement with one's own thinking; now it can be, in part, curation.

Much of human work, however, was never solitary. Transcribing, organizing and publishing the lectures of an entire conference was, until recently, a months-long project for a team of five people, with back-and-forth, corrections, discussions, lunches. It generated income and it generated companionship. Today it runs in one night, without anyone touching the material between the professor who spoke and the student who will watch. What disappears is not only the job. It is the pretext for the encounter.

The question that arises deserves to be taken seriously: if most everyday human interaction has some practical function, solving something, deciding something, teaching something, what is left when that function is handed over to the machine? A son who lives far away calls his parents to learn how to improvise dinner from whatever is in the fridge. The call lasts an hour, and the recipe takes up five minutes of it. A voice assistant would answer better and faster. If the son starts asking the machine, the parents lose a small function and, with it, the hour of conversation that came along. Multiply this by every silly question a child asks, which a pair of smart glasses will soon answer before the child even turns to their mother.

These exchanges have a socio-emotional function far greater than the information they carry. They build a repertoire of attachment that the species has always taken for granted. And it is already possible to estimate part of the effect by looking at the generation that grew up with the previous technological mediation, that of smartphones and social networks. In the United States, the proportion of 16-year-olds with a driver's license fell from 46% in 1983 to 24% in 2014, and that of 18-year-olds from 80% to 60%. In Cigna's 2020 loneliness index, with more than ten thousand respondents, 79% of Generation Z described themselves as lonely, compared with 71% of millennials and 50% of baby boomers, and heavy social media use tracked loneliness in every age group. Correlation is not causation, and there are competing economic explanations for each of these numbers. But the pattern is consistent with the hypothesis that, when technology takes over the practical function of the encounter, the encounter does not survive on its own.

There is an even deeper point, an almost hedonistic one. Whoever uses artificial intelligence to eliminate the distress of the tasks they hate also removes the contrast that gave meaning to relief. Without the low of the bad, the good stops being a peak and becomes a plateau. And plateaus always demand more. The observation is an old one in the psychology of hedonic adaptation, but it takes on another scale when the tool that removes discomfort is available for everything, all the time.

Faster than adaptation

The third concern is the one that ties the other two together: speed. We have acquired the ability to advance technology at a pace far greater than that of social and psychological adaptation. A few years ago, conversational systems seemed only slightly smarter than the ones that became famous in 2010. Today they automate entire professions, contribute to research in theoretical mathematics, and are built into glasses, phones and urban surveillance systems. The world changes quickly; the way we engage with it changes slowly, because it is made of habit, attachment and development, things that have their own pace.

The philosopher David Collingridge described in 1980 the dilemma that this asymmetry produces. When a technology is new, its effects are not yet known, and it is therefore hard to regulate; by the time the effects finally become visible, the technology is so embedded in social life that changing it costs too much. The information we would need in order to decide always arrives late. In the case of artificial intelligence, we are collectively betting that the outcome will be good. Perhaps it will be. But it is a bet, not a conclusion.

The recent history of screens shows how the dilemma plays out in practice. When home computers and later smartphones spread, psychology had no data on what they would do to child development. The most science could offer were hypotheses and precautionary recommendations. Solid evidence came later, when use was already universal. A longitudinal study of 2,441 Canadian children, published in *JAMA Pediatrics* in 2019, showed that more hours of screen time at 24 and 36 months predicted poorer performance on developmental screening tests at 36 and 60 months, and not the other way around. The effects are modest in magnitude, but the direction is clear. A smaller brain-imaging study in preschoolers, published the following year, associated screen use above pediatric recommendations with lower white matter integrity in tracts linked to language and literacy. None of this existed when parents in the 2000s were deciding whether or not to give their child a tablet. They decided in the dark, and only now do we know part of what those decisions cost.

With generative artificial intelligence, we are once again at the beginning of that curve, with one difference:

the technology mediates not only access to information and entertainment, but thinking itself and conversation itself. Children who grow up asking an assistant what they would once have asked their parents will have childhoods different from any childhood ever studied. By the time data on those childhoods exist, they will already be over.

There is a further aggravating factor that the history of screens teaches. No technology is neutral with respect to values; every technology is developed for some purpose, by someone, and the purpose of those who sell attention is to capture attention. The large platforms know precisely the average screen time of their users, how many ads fit into that time and how much each age group is worth per day. The engineers on the other side are paid to make the product as engaging as possible, and they are among the best in the world at what they do. Expecting these companies to regulate themselves, or governments that depend on them to do so in time, runs against the entire experience of the last two decades. Whoever has to set limits will once again be the individual.

What is left to the individual

If regulation will not come from outside, or will come late, the only sphere where judgment can still be exercised is the small sphere each person controls: the home, the table, the bedroom, the hours spent with loved ones. The response is modest and out of proportion to the size of the problem, but it was the same one left to parents during the explosion of screens, and those who adopted it out of precaution, before there were data, now seem to have been right. Slowing down within one's own life what cannot be slowed down in the world is a way of buying time to observe, with some density of reality, what each technology does to those who use it.

The first principle is caution, understood not as refusal but as the opposite of the automatic. No one is going to live under a rock; coexisting with these tools is inevitable and, to a large extent, desirable. The question is whether that coexistence happens on our terms or on the terms of those who profit from it. On autopilot, the other side always wins; the hand reaches for the device before the intention has formed. Consciously stopping to rethink each habit is laborious, and for that very reason it must be practiced, as one practices anything that does not come for free.

The second principle is that example weighs more than words. Children and adolescents learn their relationship with technology by watching that of adults, and a parent who preaches limits with a phone in hand teaches the phone, not the limit. It is worth remembering that the creators of these technologies have always known this: Steve Jobs said in 2010 that his children did not use the iPad and that the family limited technology at home, and the same pattern repeats among executives in the industry.

Whoever develops the product and denies it to their own children knows something the consumer does not.

In practice, this means creating a household culture with named rules. Agreeing that the time a couple spends together, outside work, is phone-free, and holding each other to it when the hand slips. Setting aside a daily period when the devices stay in another room and people lie on the sofa with a book. Deliberately replacing hours that used to go to work or to the screen with social and analog activities, even if they seem unproductive. Preferring the long video, with a beginning, middle and end, to the endless stream of cuts that leaves no time for any of the three. None of this is sophisticated. The hard part is not knowing what to do, but doing it every day against engineering designed to keep it from being done.

There are signs that this works, and that young people themselves notice when it is missing. It is not unusual for a teenager to ask their parents to reinstate a time limit that had been removed, because they cannot manage to stop on their own. A request like that is the best indicator that the work at home is being done: the child has learned to recognize capture and to name it. For the next generation, which will grow up with tools far more persuasive than ours, this capacity for self-limitation will matter more than any content the school teaches.

The same applies to artificial intelligence in intellectual work. If memory is formed through effort, one has to choose where effort is still worth it: which texts are still written by hand, which problems are still solved without help, which conversations are still held by phone even when an assistant would answer better. Delegating everything that can be delegated is rational in the short term and impoverishing in the long term. The question to ask of each task is not only whether the machine does it better, but whether it is worth having the machine do it.

Anxiety as a signal

Those who closely follow the advance of artificial intelligence tend to describe the same feeling: a frantic race to avoid being run over, while most people carry on with their lives as if nothing were happening. It is tempting to treat this distress as a problem to be solved, one more task to delegate. But anxiety is also a defense mechanism. It exists to signal danger, and there is real danger when a technology enters every home before anyone knows what it does to memory, to relationships and to childhood.

The anxiety growing among parents, physicians and educators is not a symptom of maladjustment. It is the proportionate response of those who have realized they are gambling with their own lives and, worse, with their children's lives, without the data they would like to have.

It is not wise to calm it completely. It is wise to turn it into attention: into named household rules, into effort preserved where effort still builds memory, into encounters kept up even when they have lost the practical function that justified them.

Perhaps twenty years from now we will discover that everything went well, that the species integrated one more tool as it integrated writing, the printing press and the telephone. Perhaps. Until the answer comes, the only defensible position is that of someone who uses technology with both hands and, even so, keeps a space where it does not enter. It is little. But it is what falls to each of us, and it is probably what will make a difference.

References

1. Slamecka NJ, Graf P. The generation effect: delineation of a phenomenon. *J Exp Psychol Hum Learn Mem.* 1978;4(6):592-604.
2. Bjork RA. Memory and metamemory considerations in the training of human beings. In: Metcalfe J, Shimamura A (eds). *Metacognition.* MIT Press; 1994. See also Bjork & Bjork, *Introducing desirable difficulties.*
3. Kosmyna N et al. Your Brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing task. *arXiv:2506.08872*, 2025 (preprint).
4. Sparrow B, Liu J, Wegner DM. Google effects on memory: cognitive consequences of having information at our fingertips. *Science.* 2011;333(6043):776-778. Failed replication: Camerer CF et al. *Nat Hum Behav.* 2018;2:637-644.
5. Sivak M, Schoettle B. Recent decreases in the proportion of persons with a driver's license across all age groups. *UMTRI-2016-4*, 2016.
6. Cigna. *Loneliness and the Workplace: 2020 U.S. Report.* 2020.
7. *Smithsonian Magazine.* ChatGPT helped write this award-winning Japanese novel. Jan 2024.
8. Collingridge D. *The Social Control of Technology.* Frances Pinter; 1980.
9. Madigan S et al. Association between screen time and children's performance on a developmental screening test. *JAMA Pediatr.* 2019;173(3):244-250.
10. Hutton JS et al. Associations between screen-based media use and brain white matter integrity in preschool-aged children. *JAMA Pediatr.* 2020;174(1):e193869.
11. Bilton N. Steve Jobs was a low-tech parent. *The New York Times*, Sep 2014 (reprinted by the Irish Times).