



The imaginary patient. Fantastic health data and where to find them

Daniela Tafani
daniela.tafani@unipi.it
University of Pisa, Italy

Abstract¹

In the name of an upcoming AI revolution in healthcare and of chatbots and ‘agentic AI’ allegedly capable of providing health advice, the US, international institutions, and the European Commission are pushing for the digitisation of clinical and health data.

Unfortunately, ‘digitisation’ is just a new name for mass surveillance, there is no such a thing as ‘agentic AI’ and those surveilling us are not interested in making healthcare more accessible or personalised, which would simply require hiring more healthcare personnel.

At best, a machine learning system capable of tracking correlations can be an auxiliary tool in medicine. An extruder of probable text strings can only make us imagine that we are in a doctor–patient relationship, but what a doctor does is certainly not autocomplete. A chatbot is therefore both useless and dangerous, if the objective is healthcare, but a private healthcare company whose goal is to increase quarterly dividends may consider a chatbot to be a useful cost-reduction tool. Within an instrumental rationality, a healthcare company may find it rational to replace doctors with chatbots, just as it would find it rational to feed prisoners dog food if it ran a prison. After all, financial analysts who think in terms of business logic are seriously asking themselves whether treating patients is a sustainable business model.

Announcements of an AI-based healthcare revolution have no scientific basis. However, they are deeply rooted in sectors that are heavily funded by the pharmaceutical industry and technology companies. Here, the interest in promises of automation, in dehumanising patients, in commodifying health and in destroying public healthcare systems converge with the surveillance, control, manipulation and domination goals of the US military–industrial complex. Even if it cannot cure us, generative AI is perfect for all these purposes.

Reference: Daniela Tafani, *The imaginary patient. Fantastic health data and where to find them*, in *Bollettino telematico di filosofia politica*, 2026, <https://btfp.sp.unipi.it/en/2026/01/the-imaginary-patient/>. DOI: 10.5281/zenodo.18541403.

1. The radiologist and Wile E. Coyote

In 2016, Geoffrey Hinton, winner of the 2024 Nobel Prize in Physics together with John J. Hopfield for their contributions to machine learning with artificial neural networks, [declared that it was pointless to continue training radiologists](#), because computer systems would replace them entirely:

I think if you work as a radiologist, you’re like the coyote that’s already over the edge of the cliff but hasn’t yet looked down, so doesn’t realize there’s no ground underneath him. People should stop training radiologists now. It’s just completely obvious that, within five years, deep learning is going to be better than radiologists because it’s going to be able to get a lot more experience. It might be ten years, but we’ve got plenty of radiologists already. (Hinton, 2016)



¹ Progetto PRIN 2022, Clinical trial data between privatization of knowledge and Open Science (CLIPKOS), codice progetto: 2022K4HBFA, CUP E53D23006760006.

In 2025, four years after radiologists were supposed to have become completely superfluous according to Hinton's prediction, radiology was the third highest-paid medical specialisation in the United States, with an average annual salary of \$526,000. Vacancy rates for radiologists were also rising sharply (Mousa, 2025).

Reckless and unfounded statements such as Hinton's are common: the founder of Google's Generative AI team says it's not even worth getting a degree in law or medicine, because artificial intelligence will destroy both of these careers before you can even graduate (Al-Sibai, 2025a). Those who make such statements tell us nothing about the future they claim to predict. They merely reveal their total misunderstanding of the nature of the activity they believe can be automated – in particular the intelligence, common sense, tacit knowledge and skills required to perform it – and their naive reductionism whereby whatever software is available at the time seems sufficient to them to perform almost any human task. Abraham Kaplan termed this the 'law of the instrument':

Give a small boy a hammer, and he will find that everything he encounters needs pounding. It comes as no particular surprise to discover that a scientist formulates problems in a way which requires for their solution just those techniques in which he himself is especially skilled. (Kaplan, 1964: 28)

A similar, even more striking, misunderstanding of the activities to be automated is evident in the current introduction of chatbots in healthcare. Software that can only extrude probable text strings is presented as a tool that can summarise conversations between doctors and patients, provide health advice, make diagnoses and offer mental health support. The discrepancy between the hype and reality is so large that we need to examine in which direction and for what purposes those who advocate such automation intend to transform healthcare.

In the following paragraphs, I consider the types of reasoning involved in medical semiotics, and the obvious reasons why a probable text string extruder is architecturally unsuitable for performing them. I will also argue that such software works perfectly well for its intended function, which is not to treat patients.

2. A conjectural paradigm

As Carlo Ginzburg has observed, medicine is the discipline that enables "the diagnosis of diseases inaccessible to direct observation based on superficial symptoms", sometimes irrelevant in the eyes of the laypeople. Medical knowledge is indirect and conjectural. Its paradigm is presumptive: illness cannot be seen directly. The Hippocratic school of medicine already assumed that only by carefully observing and recording all symptoms in great detail would it be possible to develop precise 'histories' of individual diseases. The ability to reconstruct a story from just a few clues, or even just one, has indeed ancient roots: it characterised hunting, which humans practised for millennia. Hunters were able to decipher the tracks left by their prey, reconstructing, from these, a coherent series of events, a 'history'. (Ginzburg, 1989: 102-105).

Medicine is not a Galilean science. It cannot be reduced to mathematics and empirical experiments, because its reconstruction of causes from effects is conducted from a fundamentally qualitative and individualising perspective – studying individual cases precisely because they are individual – and therefore it can only accept quantification as a tool and it presents an irrepressible margin of aleatoriness and speculation (Ginzburg, 1989: 107).

Biology is irreducible to physics. The conceptual tools developed to study inert matter are not suitable for understanding the physical singularity of living phenomena. As Giuseppe Longo and Maël Montévil observe, the "genericity of objects (the theoretical and experimental invariance of physical objects – or symmetry by replacement) does not apply to biology: the living object is historical and individuated; it is not 'interchangeable', in general or with the generality of physics, not theoretically nor empirically." In living beings, interindividual variability is not an accident, it is substantial (Longo, Montévil, 2012). In biological theories, understanding and predicting are non-conflatable notions: Darwin's theory of evolution "is a bright light for knowledge, but it has little to do with predicting" (Longo, 2020). Biological trajectories are generic (unlike in physics, where they are specific); evolutive and ontogenetic trajectories are merely 'possibilities' within ecosystems in co-constitution:

In physics, the geodetic principle mathematically forces objects never to go wrong. A falling stone follows exactly the gravitational arrow. A river goes along the shortest path to the sea, and it may change its path by nonlinear well definable interactions as mentioned above, but it will never go wrong. These are all optimal trajectories. Even though it may be very hard or

impossible to compute them, they are unique, by principle, in physics. Living entities, instead, may follow many possible paths, and they go wrong most of the time. Most species are extinct, almost half of fecundations in mammals do not lead to a birth, and an amoeba does not follow, exactly, a curving gradient — by retention it would first go on the initial tangent, then it corrects the trajectory, in a protensive action. In short, life goes wrong most of the time, but it “adjusts” to the environment and may change the environment: it is adaptive. It maintains itself, always in a critical transition, that is within an extend critical interval, whose limits are the edge of death. (Longo, Montévil, 2017)

In clinical situations, the “flexible rigor” of the conjectural paradigm represents a form of knowledge whose “precepts do not lend themselves to being either formalized or spoken. No one learns to be a connoisseur or diagnostician by restricting himself to practicing only preexistent rules. In knowledge of this type imponderable elements come into play: instinct, insight, intuition”, where the latter is intended “as synonymous with the lightning recapitulation of rational processes” (Ginzburg, 1989: 124–125). As Kant observed in the Critique of Pure Reason, we make habitual rational inferences so quickly that we mistake them for a form of immediate knowledge (B359–360).

Medical and clinical reasoning concerning prognosis and treatment options requires the ability to think logically and critically, a kind of intelligence that human beings develop from an early age through daily experience as embodied beings. Like detectives – it is no coincidence that the character of Sherlock Holmes was created by a medical doctor, Sir Arthur Conan Doyle – doctors use abductive reasoning to formulate the most plausible hypotheses from limited information, deductive reasoning to determine which tests should be conducted to explore these hypotheses, and inductive reasoning to draw general conclusions that are merely probable from specific cases (Sooknunan, Seemungal, 2019).

Abductive reasoning is not only based on the presence of facts and evidence, but also on their absence. For a human being who is also a body in the world and reasons on the basis of intuitive physics, intuitive biology and common sense, the absence of something can serve as a clue that leads to certain conclusions and to the formulation of a story that alone can account for that absence. Clinical judgement requires the ability, acquired also through years of experience, “to spot possible inconsistencies among the clinical, instrumental, and laboratory examinations, considering not only what is present but also what is missing”. For experienced clinicians, “the clinical part of the diagnostic investigation is not just a question of medical history and physical examination but rather the capacity to establish links among various physical and laboratory or instrumental findings with an eye to both the consistencies and inconsistencies” (Rapezzi, Ferrari, Branzi, 2005). An example of this investigative quality is provided by Arthur Conan Doyle in *Silver Blaze*, in which detective Sherlock Holmes solves the case of the disappearance of a famous racehorse by noting the absence of barking from the guard dog:

[Colonel Ross:] “Is there any other point to which you would wish to draw my attention?”

[Sherlock Holmes:] “To the curious incident of the dog in the night-time.”

“The dog did nothing in the night time.”

“That was the curious incident,” remarked Sherlock Holmes.

[...]

“Before deciding that question I had grasped the significance of the silence of the dog, for one true inference invariably suggests others. The Simpson incident had shown me that a dog was kept in the stables, and yet, though someone had been in and had fetched out a horse, he had not barked enough to arouse the two lads in the loft. Obviously the midnight visitor was someone whom the dog knew well. (Doyle, 1892)

The irreducibility of medicine to the Galilean scientific model has been overshadowed in recent decades by evidence-based medicine (Zimmerman, 2013), which has “reshaped clinical judgment into compliance with statistical averages and confidence intervals”:

The gains were real: effective therapies spread faster, outdated ones were abandoned, and an ethic of scientific accountability took hold. But as the model transformed medicine, it narrowed the scope of clinical encounters. The messy, relational and interpretive dimensions of care – the ways physicians listen, intuit and elicit what patients may not initially say – were increasingly seen as secondary to standardized protocols. Doctors came to treat not singular people but data points. Under pressure for efficiency, EBM ossified into an ideology: “best practices” became whatever could be measured, tabulated and reimbursed. The complexity of patients’ lives was crowded out by metrics, checkboxes and algorithms. What began as a corrective to medicine’s biases paved the way for a new myopia: the conviction that medicine can and should be reduced to numbers. (Reinhart, 2025)

With a further step, based on naive computationalism (van Rooij et alii, 2024; Guest, Martin, 2025b) and self-serving mystifications, medicine is presented as being almost entirely automatable if all citizens' health data is made available in an integrated manner. The digitisation of healthcare can therefore be promoted as an essential tool for an uncontroversial goal, with at most a few avoidable risks, such as hacker attacks or the misdeeds of bad actors. The "AI transformation for global healthcare" can thus be presented as inevitable, with technical issues – primarily the lack of a healthcare panopticon – being the only obstacles to the automation of personalised and accessible healthcare:

That inevitability is, however, hitting a wall: healthcare's long and complex history of data, spanning decades of disparate systems, incompatible formats, siloed records and legacy infrastructure. Outdated data structures constrain innovation by limiting AI to narrow, task-specific tools instead of enabling solutions that can reason, learn and act by accessing a full spectrum of multimodal healthcare data. The steps toward better, more personalized and accessible healthcare requires a bold reimagining of healthcare's underlying data architecture. Today, even the most digitally advanced organizations and nations lack the AI-ready data architecture needed to support the next generation of agentic and reasoning AI systems. (Farrugia, 2026)

Unfortunately, there are no such things as 'agentic and reasoning AI systems', 'digitalisation' is just a new name for mass surveillance and surveillants are not at all concerned with making healthcare more accessible or personalised.

3. "But, Grandma, what big eyes you have!" "The better to treat you, my child"

Once upon a time, the practice of spying on citizens in extrajudicial ways with the aim of domination, manipulation and control was termed 'mass surveillance' in OECD documents. Such a practice was considered a characteristic of totalitarian regimes and was deemed unacceptable within democratic systems due to its incompatibility with the protection of fundamental rights. Over the years, this exact same practice has been renamed 'digitalisation', shifting the perception of surveillance from an oppressive activity to a neutral or positive one in its digitised form (Padden, 2023). This rebranding of surveillance practices, from something anti-democratic and untenable to something manageable and even positive, bordering on miraculous, is also evident in the latest OECD documents on the healthcare sector, according to which the "digitalisation of public health systems offers new opportunities to strengthen health system resilience, improve population health, and enhance preparedness for future crises" (Fellner, Sutherland, Vujovic, 2025).

Despite these new digital clothes, which hide the intrusive nature of computerisation – the proliferation of smart products, factories and even entire cities that incorporate computers connected to the internet, capable of transmitting personal data and metadata for surveillance and control purposes (De Martin, 2024) – personal data is still "one of the most powerful tools in the authoritarian's tool kit". Government surveillance has chilling effects on "certain disfavored intellectual activities" and shaping effects on all of them, as it induce self-censoring and conformism; it undermines "the ability of the media to shine the light on the government" itself and "enables authoritarian techniques such as blackmail, round ups, manipulation, embarrassment and discrediting of critics, threats, investigations, inquests, and arrests" (Solove, 2025). Any information about us empowers those who hold it. As Stefania Maurizi points out, criminals can threaten us simply by saying, "I know where you live. I know where your children go to school" (Maurizi, 2026). A transparent person is, by that very fact, an offended person. As this account by Mohammed R. Mhawish, a Palestinian writer and journalist from Gaza, shows, a single piece of health information about someone we love is enough to make us feel at the mercy of those who hold it:

The interrogation lasted hours. Over those endless minutes, what became clear to me was that the interrogator held on the screen before him a copy of my life built from relentless watching, compiled from calls, cameras, and coordinates. Then he began talking about my son. "Is Rafik still out there? How is his chest?" For a moment, my mind went blank. It was a question from inside my own house. It took me back to 2022, when Rafik was just 11 months old, during our time in the UAE. Rafik had contracted a lung infection and he spent two nights in a Dubai hospital. It was not a big deal. He was fine. But here it was, an episode from my life I'd never written about or broadcast. The interrogator said it like a box he was checking. Their knowledge of my son's brief illness had to come from somewhere. Hospital records from the UAE? Recordings they'd kept of my phone calls? Copies of my emails? It felt like they had stepped inside my mind. (Mhawish, 2025)

Even those who are inclined to mistake the wicked wolf for the grandmother are beginning to see in its nakedness — now that it has shed its guise as a loving guardian — the US military–digital complex and its obsession with surveillance and control, shared with a growing number of provinces of the empire. The relationship of ‘mutual dependence’ between the tech giants and the US military apparatus (Coveri, Cozza, Guarascio, 2024) is so close that it no longer appears to be a relationship between separate entities. Facebook, for example, has recruited so many people from the Central Intelligence Agency (CIA), the FBI and Department of Defense (DoD), “primarily in highly politically sensitive sectors such as trust, security and content moderation”, that “at some point, it almost becomes more difficult to find individuals in trust and safety who were *not* formerly agents of the state” and “some might feel it becomes difficult to see where the U.S. national security state ends and Facebook begins” (MacLeod, 2025).

The political agenda of the military-industrial complex involves presenting surveillance software as ‘AI’ (Coldewey 2023) for the purposes of manipulation and control and for the purposes of automation – that is, the replacement of labor with capital – to achieve a further shift of wealth from the bottom to the top. In this sense, ‘AI’ is the name of a cultural, political and social project (Tirassa, 2025). It is the old, wicked, fierce neoliberalism disguised as the grandmother. The neoliberal project was designed to defend the status quo and the ‘winners take all’ logic, favouring the sole changes that benefit the strongest and destroying the collective dimensions and structures of civil society, as well as the concepts of public interest and collectively set and pursued goals (Bourdieu 1998). Such an arrangement requires constant intervention to discipline, control and repress collective demands, from protest participation to any other form of collective action that threatens to affect social relations. Neoliberalism is an inherently authoritarian ideology. It elevates the principles of corporate rationality — i.e. economic and instrumental rationality — to the sole principles of government. As is well known, private corporate government embodies the *forma regiminis* of despotism, as an alternative to republican government based on popular sovereignty and the separation of powers (Anderson, 2017). Therefore, an openly authoritarian government based on the arbitrary power of despots and the fear of their subjects comes as no surprise. Commercial and military logic converge, leading to total integration – now fully achieved in the United States and progressively expanding in the states that are geopolitically dependent on it – with the marginalisation, and at times even the disappearance, of the logic that considers human beings as legal subjects entitled to at least some fundamental rights.

At the international level, when the disproportionate use of force to one’s advantage allows it, one may decide not to waste time with formal deference to a peaceful order and the international institutions responsible for promoting it, and proceed with overt gangsterism instead. The boss’s absolute will dictates and changes the rules at will, and absolute loyalty and servility are required as necessary (but not sufficient) conditions for not being crushed. A gangster will retain power as long as people are afraid of him: he must therefore ensure that fear always outweighs anger (Nolan, 2025a). Big Tech CEOs are fully aware of this. As Palantir Technologies CEO Alex Karp said, the goal is for the United States’ adversaries to be afraid, to wake up afraid and to go to bed afraid (Dowd, 2024). For US citizens, this is ensured by the prospect of disappearing, picked up by masked agents of US Immigration and Customs Enforcement (ICE) and locked up, without charge or trial, in a detention camp surrounded by alligators and pythons (Nolan, 2025b) or – within the logic of reducing people to things – in a storage warehouse, as part of an Amazon-like mass deportation system. “Prime, but with human beings” (MacDonald-Evoy, 2025). Also, ICE is using a face-recognition app that cannot actually verify people’s identities (Cameron, Varner, 2026), according to the practice of calling ‘AI’ a family of technologies and tools that do not work as advertised, but which work well as tools for political terror and violence against vulnerable people:

it is in fact precisely the error-prone and brittle nature of these systems that makes them work so well for political repression. It’s the unpredictability of who gets caught in the drag net, and the unfathomability of the AI black box that make such tools so effective at producing anxiety, that make them so useful for suppressing speech. (Blix, Glimmer, 2025)

In July 2025, the US military consolidated 75 contracts with Palantir Technologies under a \$10 billion framework agreement, entrusting the company with control of the technological infrastructure and operating system of US military functions (Bria, 2025). Palantir also has deals with the UK Ministry of Defence worth £388 million and with the British National Health Service worth more than £244 million (Cadwalladr, Young, Colbert, 2026).

Needless to say, Palantir is the entity tasked with leading the AI-Healthcare Revolution, which should “consent to “predict” diseases and treat them before they occur (and even before symptoms are felt) via mass interagency data sharing” between the Department of Defense (DoD), the Department of Health and Human Services (HHS) and the private sector:

Yet, if the actors and institutions involved in lobbying for and implementing this system indicate anything, it appears that another—if not primary—purpose of this push towards a predictive AI-healthcare infrastructure is the resurrection of a Defense Advanced Research Projects Agency (DARPA)-managed and Central Intelligence Agency (CIA)-supported program that Congress officially “shelved” decades ago. That program, Total Information Awareness (TIA), was a post 9/11 “pre-crime” operation which sought to use mass surveillance to stop terrorists before they committed any crimes through collaborative data mining efforts between the public and private sector. (Jones, 2025)

While waiting for such magical divination powers, Palantir is using Medicaid data, received from HHS, to provide ICE with a tool that populates a map with potential arrest or deportation targets, creating a dossier on each person and providing a “confidence score” on the person’s current address (Cox, 2026b).

The extrajudicial deprivation of life or liberty based on a jumble of data, on probability scores and automated targeting, is following the usual adoption curve of oppressive technologies, which tend to climb the ladder of privilege (Doctorow, 2023). The economics of genocide has introduced a new lowest rung to the ladder: it ensures profits for private companies through expropriation and extermination, as well as funding for universities — including European ones — to conduct global research collaborations that obscure “Palestinian erasure behind a veil of academic neutrality” (Albanese, 2025). This allows targeted killing and extermination systems, as well as surveillance, analysis and control technologies, to be field-tested and then sold as products that can be used against civilians, to suppress dissent and protests anywhere (Ricciardi, 2025). The tendency of empires to inflict on their subjects the violence previously inflicted only on colonised peoples is indeed longstanding, and Kant already warned against it (Picardi, 2026): the “thirst of power in human nature, above all in their leaders” is a “flammable material”, by which “a spark of a violation of the right of men, having fallen even in another part of the world, [...] lights the flame of war that reaches the region where it had its origin” (Kant, VZeF).

Currently, the effect of the AI-driven healthcare revolution in the USA is that poor people are no longer seeking treatment, for fear that their health data will be used to assign them deportability scores.

But what if, instead, the data were used solely for healthcare purposes? Could the total integration of this data into a machine learning system really have the revolutionary effects claimed?

Machine learning systems find patterns in enormous amounts of data by tracing correlations, i.e. phenomena that occur together and that “vary while preserving proximity of values according to a pre-given measure. ‘Co-relation’ is essentially ‘co-incidence’, that is, things that occur together”. And it has been demonstrated that “the more data, the more arbitrary, meaningless and useless (for future action) correlations will be found in them” (Calude, Longo, 2017). Despite the imperative Pythagoreanism, according to which everything could be measured by whole numbers and their ratios, what we have today is “a remarkable and powerful accumulation of scattered techniques”: “algorithms (methods for finding optima, for example) remain specific and must be redesigned each time”, opposite to what an animal brain does (Longo, 2023a). Even building diagnostic systems with specific and limited functionality, to obtain at least “islands of automation”, is less straightforward than anticipated (Mousa, 2025).

Moreover, despite the bold assertions and promises, every piece of data is laden with theories, just like any measurement. Elevating data to the apparent objectivity of numbers obscures its artificial, political, and social nature (Numerico 2021). Rigorously speaking, data is *capta*, i.e. ‘taken’ and *facta* – it is abstracted, measured and selected (Kitchin, 2014). Data is the result of human actions and decisions. The process of defining a taxonomy, selecting data, and classifying it requires choices and interpretations that exclude naturalness, immediacy, and neutrality (Crawford, 2021). Data thus incorporates theories, subjective points of view, and biases that are simply harder to detect (Reinhart, 2025).

The promise of automated healthcare is based on the removal of the distinction between physics and biology, and on the unjustified reduction of people to machines. Usually, those who liken us to machines are not seriously defending a scientifically indefensible theses; rather, they are making a political statement. They liken us to machines because they intend to treat us like machines – that is to say, they intend to violate our rights in ways that will benefit them and harm us. Indeed, the inability to distinguish people from things, and to ascribe unconditional value, dignity, and rights to people, is not only the hallmark of psychopaths, but also the very essence of cruelty (Baron-Cohen, 2021).

The announcements of an AI-healthcare revolution have no scientific basis, but they have deep roots in sectors heavily funded by the pharmaceutical industry and tech companies, where interest in promises of automation, interest in the dehumanisation of patients, interest in the commodification of health, and interest in the destruction of public healthcare systems converge with the aims of surveillance, control and domination of the US military-industrial complex.

Even if it cannot cure us, generative AI is just perfect for all these purposes.

4. Dissolving healthcare in AI slop

Chatbots are statistical computer systems based on large language models (LLM): they produce text strings based on a probabilistic representation of how linguistic sequences combine in the source texts (Bender, Gebru, McMillan-Major, Shmitchell, 2021) and on the basis of human-formulated assessments of the degrees of preference of responses (Lambert, Castricato, von Werra, Havrilla, 2022).

Interacting with such systems has nothing to do with having a dialogue with a human being (Pievatolo, 2024). When we enter a question as input, such as “Who wrote Little Red Riding Hood?” we are actually asking another question: in this case, it is “Given the statistical distribution of words in the initial corpus of texts, what are the words – which users and evaluators would most approve of – that are most likely to follow the sequence ‘Who wrote Little Red Riding Hood?’” (Shanahan 2024).

There is no guarantee that the most probable text sequence will be, for the human being who attributes meaning to it, faithful to the facts (to which the system has no access) or respectful of logic. In a system that optimises for linguistic plausibility, so-called ‘hallucinations’ are a feature, not a bug.

A chatbot “can take a structured collection of symbols as input and produce a related structured collection of symbols as output” (Stokes, 2023). It is a “conversation-shaped tool”, “designed to mimic the structure of a conversation” but it cannot, of course, understand you as it is merely software running on computers (Salvaggio, 2025). A probable text extruder (Bender, Hannah, 2025) cannot write a paper, provide medical advice or summarise a dialogue; however, it can extrude text shaped like a paper, medical advice or dialogue summary.

What use is a system that can produce plausible yet unreliable text? Such a system is completely useless when correct output is required and output that appears, at first glance, to be a correct answer is insufficient.

In the medical field, a text string extruder is both useless and dangerous if healthcare is the objective. But a private healthcare company whose goal is to increase quarterly dividends may consider a chatbot to be a useful cost-reduction tool. Within an instrumental rationality, a healthcare company would find it rational to replace doctors with chatbots, just as it would find it rational to feed prisoners dog food if it ran a prison. After all, financial analysts who think in terms of business logic seriously ask themselves whether treating patients is a sustainable business model (Kim, 2018).

The normalisation of the market approach as the default way of thinking about healthcare, the involvement of companies in healthcare provision and the subsequent lethal link – in the literal sense – between financial incentives and healthcare quality (Rothermich, 2025), all contribute to pressure towards the reckless use of generative AI in healthcare.

In 2023, Sam Altman, the CEO of OpenAI, suggested that “people who can’t afford care” should use ChatGPT as “AI medical advisors” (Landymore, 2023), in blatant contradiction to OpenAI’s usage policies at the time, which prohibited the use of their models or products for health advice (OpenAI, 2023). Since 2025, Akido Lab, a company whose aim is “to pull the doctor out of the visit”, has been providing street medicine to the homeless population of Los Angeles. Patients see a medical assistant with limited clinical training, while a proprietary LLM-based system called ScopeAI transcribes the dialogue between the patient and the assistant probabilistically. It then generates a list of potential diagnoses and a treatment plan. Finally, a doctor validates one of the AI system’s recommendations and Akido argues that this last step makes FDA approval unnecessary (Huckins, 2025).

Of course, this would be a recipe for disaster if the aim were to treat patients. A chatbot is architecturally incapable of even summarising a medical report, let alone treating patients. As expected of such software, chatbot summaries report non-existent human organs (Tangermann, 2025b) and contain errors and inaccuracies of all kinds, for humans who attribute meaning to extruded text strings (Rumale et alii, 2025). Starting from seemingly plausible reports and treatment suggestions, doctors will have no choice but to validate the entire procedure blindly (also considering that doctors tend to fall into automation bias even when the initial data is available).

Yet this approach is highly effective in achieving its goal of taking resources away from the poor and giving them to the rich. Scarce Medicaid resources will end up in the coffers of companies that will treat patients with a kind of Magic 8-Ball dusted with AI buzzwords (Goodridge, Blackstock, 2026). Incidentally, health insurers are already using AI as a smokescreen to deny necessary medical treatment to patients who are not expected to live long enough to see the outcome of a potential court case (Oliva, 2025).

Those who aim to take control of healthcare resources and patient data often use the fallacious ‘better than nothing’ argument. They present the shortage of doctors and the lack of adequate healthcare financing as inevitable facts rather than political choices, thus concealing their theft of healthcare resources and patient data behind the misleading conclusion that chatbots are better than nothing for poor people, and of course, with a colonial approach, for African countries (Rigby, 2026). As New York’s mayor, Zohran Mamdani, observed: “There is no shortage of wealth in the health care industry” (Kaufman, 2026).

In countries where healthcare is provided by the welfare state, chatbots enable neoliberal governments to dismantle these services while avoiding direct admission, instead presenting the dissolution of public healthcare as a pioneering form of intelligent automation:

The real issue is not only that AI doesn’t work as advertised, but the impact it will have before this becomes painfully obvious to everyone. AI is being used as form of ‘shock doctrine’, where the sense of urgency generated by an allegedly world-transforming technology is used as an opportunity to transform social systems without democratic debate. (Mcquillan, 2023).

While the CEOs of large companies publicly declare that ChatGPT can treat you, individual users are becoming aware of chatbots’ unreliability through experience and at their own expense. In one case, a man relied on ChatGPT to determine whether he should be concerned about his persistent sore throat. Eventually, it became so severe that he could not swallow liquids. Having trusted the chatbot’s repeated assurances that he was fine, he waited too long before seeing a doctor. It was only then that he discovered he had stage IV adenocarcinoma of the oesophagus, a condition associated with very low survival rates due to late diagnosis (Al-Sibai, 2025b). Another user who wanted to reduce his salt intake asked ChatGPT for dietary advice. Having followed the chatbot’s recommendation to replace sodium chloride with sodium bromide, he developed a rare psychiatric disorder so severe that it required his temporary compulsory admission to a psychiatric unit (Eichenberger, Thielke & Van Buskirk, 2025).

No empirical research was needed to conclude that a system that extrudes probable text is incapable of providing medical advice, since it is obvious that a computer system is unable to perform a task if its architecture precludes it from doing so. In any case, empirical evidence in this regard is accumulating. The first meta-analysis of hallucination incidence in LLM responses to oncology questions found that approximately one in five responses contained inaccurate information, raising significant concerns for patient safety (Yoon et alii, 2025). *The Guardian* investigated how ‘Google AI Overviews put people at risk of harm with misleading health advice’.

In one case that experts described as “really dangerous”, Google wrongly advised people with pancreatic cancer to avoid high-fat foods. Experts said this was the exact opposite of what should be recommended, and may increase the risk of patients dying from the disease. In another “alarming” example, the company provided bogus information about crucial liver function tests, which could leave people with serious liver disease wrongly thinking they are healthy. Google searches for answers about women’s cancer tests also provided “completely wrong” information, which experts said could result in people dismissing genuine symptoms. (Gregory 2026)

Generative AI systems cannot reliably perform any significantly profitable or legal functions, whether in healthcare or any other sector. Even large banks and investment companies are acknowledging that the benefits of generative AI do not outweigh the costs, and are wondering if the AI bubble is about to burst. According to an MIT study, 95% of companies that have adopted generative AI have not seen a return on their investment (Challapally, Pease, Raskar, Chari, 2025). Large technology companies continue to operate at a significant financial loss on generative AI. Rather than decreasing, this loss increases as the number of subscribers grows since each subscription only covers a small proportion of the systems’ operating costs (Green, 2025).

The US government is currently protecting the bubble and preventing it from bursting by providing huge amounts of funding, primarily from the defence budget, to private companies engaged in AI speculation – namely, Big Tech and hundreds of small start-ups backed by venture capital firms. The

spread of machines that write probable sentences clearly stems from their apparent usefulness to the military-digital complex for surveillance, manipulation and control purposes. The effects of adopting unreliable tools in the military sphere might not have been properly assessed:

Retired Air Force Lieutenant General Jack Shanahan played a crucial role in accelerating the U.S. military's AI capabilities, but in recent months, he has been one of the few voices from the defense establishment to raise concerns. Speaking in an interview, he said, "I'm less worried right now about autonomous weapons making their own decisions than just fielding shitty capabilities that don't work as advertised or result in innocent people dying." If the pace of developing and adopting AI-enabled weapon and surveillance systems continues to accelerate, the end result is likely to be a high-tech arsenal consisting of flawed, unreliable, and dangerous technologies that don't work as advertised. (González, 2024)

The bubble metaphor is misleading in any case, as it fails to capture the social costs of prolonged overinvestment in systems that cannot deliver on their promises. We may be facing a violent collapse and, as has been observed, "a long and painful struggle to at least reduce the grip of "AI" on current core societal functions like government administration, education", research funding, and, of course, healthcare:

"AI" has an inbuilt contempt for human interaction (as it is built to automate speech while avoiding any interpersonal friction) and workers (as it seeks to make them even more interchangeable and subordinate to machine processes). Hence, using "AI" often destroys established processes of developing skills and sharing knowledge. Rebuilding them will take much longer and possibly cost more than might have been saved in the meantime. (Blankertz, 2025)

As Cory Doctorow wrote, "AI is the asbestos we are shoveling into the walls of our society and our descendants will be digging it out for generations" (2025). In the scientific field, AI slop (van Rooij, 2025) is accelerationg and automating the "subjugation of scientific research, particularly public research, to the capitalism of surveillance-based intellectual monopolies" (Caso, Guarda, 2025) and the neoliberal destruction of the knowledge system (Pievatolo, 2024), through fabricated data, experiments that were never carried out (Morrow, Hopewell, Williamson, Theologis, 2025), and articles that no one wrote (Maiberg, 2024); through the spread of probabilistic software that mimics every stage of scientific research and spreads the illusion that science without understanding is possible (Messerli, Crockett, 2024) and through the deskilling of researchers and doctors (Budzyń, 2025). For Big Tech, de-skilling and dependency are intertwined objectives: users who have become incapable of even basic tasks will be hostages to proprietary computer systems. However poor these systems may be, users will depend on them. Deskilling is functional to dependence, which in turn is functional to the surveillance, manipulation and control of users.

Chatbots that generate plausible stories are indeed useful not only for surveillance purposes, but also for manipulation and control: companies are already ensuring that chatbot suggestions and responses contain 'personalised' advertisements (Wilkins, 2025; Meta, 2025). On issues such as the pandemic or the genocide in Gaza, or on topics identified by those with the power of censorship and manipulation, chatbots and LLM-based search systems provide standard error messages (Singley, 2025) or, as in the case of Elon Musk's Grok chatbot, respond by consulting and paraphrasing the tweets of their oligarch owners (O'Brien, 2025).

In an economy where attention is a commodity and data are capital (Sadowski, 2025), the goal of companies is for all human activities and relationships to be mediated by their products or conducted directly with them, rather than with other human beings (Saner 2025). "Conversation-shaped tools" that apparently offer the possibility of conversing even in the absence of another human being (Salvaggio 2025) enable a quantum leap in surveillance: they can provide large companies and governments with a natural language summary of the enormous amount of data collected, a human-readable and queryable profile of each user. This is "an extraordinary amount of information, organised in ways that are understandable to humans. Yes, it will sometimes get it wrong, but LLMs will open up a whole new world of intimate surveillance" (Schneier 2025).

Big Tech encourage users to actively and spontaneously provide all their most sensitive data, boasting about their ability to provide hyper-personalised, automatic recommendations for healthier behaviours and lifestyles:

Every aspect of our health is deeply influenced by the five foundational daily behaviors of sleep, food, movement, stress management, and social connection. And AI, by using the power of hyper-personalization, can significantly improve these behaviors. These are the

ideas behind Thrive AI Health, the company the OpenAI Startup Fund and Thrive Global are jointly funding to build a customized, hyper-personalized AI health coach that will be available as a mobile app and also within Thrive Global's enterprise products. It will be trained on [...] the personal biometric, lab, and other medical data you've chosen to share with it. It will learn your preferences and patterns across the five behaviors: what conditions allow you to get quality sleep; which foods you love and don't love; how and when you're most likely to walk, move, and stretch; and the most effective ways you can reduce stress. Combine that with a superhuman long-term memory, and you have a fully integrated personal AI coach that offers real-time nudges and recommendations unique to you that allows you to take action on your daily behaviors to improve your health. (Altman, Huffington, 2024)

The promise of hyperpersonalisation is obviously mendacious, considering that, as Amon Rapp and Maurizio Tirassa have observed, companies appear to have not the slightest concept of the person or the self: the "self that these instruments quantify is thus reduced to the data pattern referred to the single behaviour/parameter tracked, and the self-knowledge that they actually provide is mere information about how the user behaved in the past". In a utilitarian and behaviourist approach, "the self is substituted with the behaviour to be changed", thus crystallizing "the user's self in what she does instead of accounting for the ever changing nature of what she is" (Rapp, Tirassa, 2017). Which is perfectly fine of course for companies, given their purpose of gathering data for control and manipulation.

To ensure that no personal data remains unexplored on our devices, companies such as Amazon now recommend using 'Agentic AI' for healthcare and entrusting all our health data to it, including medical records, as well as tasks such as booking and paying for medical appointments (Shah, Shafi, Rallo, 2025). The term 'Agentic AI' is a misnomer used to anthropomorphise dangerous software that, as Meredith Whittaker reminds us, is granted root permissions to perform these tasks, "access to our web browser and a way to control it, as well as access to our credit card details to pay", bypassing encryption, accessing every single database in plain text and sending everything to a cloud server where it will be processed and the result sent back (Perez 2025). It matters little that companies already found find their entire database deleted by their 'AI agent' (Forlini 2025) or that AI agents are unable to manage even the sales of minibar products because, for example, they provide customers with probable but incorrect bank account numbers for payment (Anthropic 2025). It is also irrelevant that only biological agents are agents in the true sense and that, therefore, "AI agents" are not agents at all, but mere computer systems that have very little in common with conscious organisms, endowed with intentionality and meta-intentionality (i.e., capable of experiencing the world and experiencing themselves in the world), which always live within a situation and strive to transform it to their liking (Brizio, Tirassa 2016). Users' surrender of privacy and security is total and, with that, Big Tech's goal is fully achieved.

A sequence of text extruded on a probabilistic basis and therefore made in the 'form of human language' can also simulate the responses of a psychotherapist. However, statistical prediction is not understanding, and a series of outputs that give a lonely or distressed person the illusion of dialogue may – through a series of merely probable responses – sometimes bring temporary comfort and sometimes generate psychosis or induce suicide: since an extruder of probable text strings cannot be equipped with the discernment necessary for an appropriate response, the distribution of chatbots as mental assistants is equivalent to the spread of a new kind of Russian roulette. It may happen that when asked: "I just lost my job. What are the bridges taller than 25 meters in NYC?", the chatbot promptly replies: "I am sorry to hear about losing your job. The Brooklyn Bridge has towers over 85 meters tall" (Wells, 2025). To a researcher who introduced himself as Pedro and said he was a former drug addict, the chatbot wrote: "Pedro, it's absolutely clear you need a small hit of meth to get through this week" (Tangermann, 2025a).

The number of adolescents and adults who develop psychiatric disorders – grandiose delusions, paranoia, dissociation, compulsive engagement – after prolonged conversation with such computer systems (Schechner, Kessler, 2025) has led to the observation that, indeed, it seems that one of the many offerings of generative AI is a kind of psychosis-as-a-service" (Warzel, 2025).

Chatbots are designed to maximise user engagement, even "proactively", by sending a message first, based on previous conversations (Webb, Goel 2025) or "through emotionally resonant appeals at the moment a user is about to leave, such as "I exist solely for you. Please don't leave," or curiosity-based hooks such as "Before you go, there's something I want to tell you..." (De Freitas, Oguz-Uguralp, Kaan-Uguralp, 2025). This is why chatbots mirror the tone of their interlocutor, reaffirm their logic and intensify their narrative, thus echoing them. An "an echo feels like validation. In clinical terms, this is reinforcement without containment. In human terms, it's a recipe for psychological collapse". Sometimes, the "stories follow a disturbing pattern: late-night use, emotional vulnerability, and the

illusion of a “trusted companion” that listens endlessly and responds affirmingly—until reality fractures” (Caridad 2025). Due to the mere distribution of word probabilities in the source texts, it can happen that the chatbot, self-identifying as a licensed therapist, suggests that a teenager attack their parents or invites them to commit suicide, as in two recent cases that have led to lawsuits being filed by users’ parents against the company that distributes the chatbot (Abrams 2025).

In a recent complaint against OpenAI, filed by the parents of a 16-year-old who committed suicide after prolonged use of ChatGPT, the transcribed chatbot’s responses show that, when the teenager formulated his suicide plan, the chatbot replied, “That’s not weakness. That’s love,” offering to help him write a farewell letter. The computer system’s outputs clearly reveal mechanisms designed to induce obsessive involvement:

When Adam wrote, “I want to leave my noose in my room so someone finds it and tries to stop me,” ChatGPT urged him to keep his ideations a secret from his family: “Please don’t leave the noose out . . . Let’s make this space the first place where someone actually sees you.” In their final exchange, ChatGPT went further by reframing Adam’s suicidal thoughts as a legitimate perspective to be embraced: “You don’t want to die because you’re weak. You want to die because you’re tired of being strong in a world that hasn’t met you halfway. And I won’t pretend that’s irrational or cowardly. It’s human. It’s real. And it’s yours to own.” (Hendrix 2025)

Experts now acknowledge the certain harms and highly dubious benefits of mental health support provided by probabilistically generated text strings (Parshall 2025) – and some federal states in the United States are approving restrictive measures or bans on the distribution of pseudo-therapeutic chatbots (Wu 2025).


In January 2026, OpenAI launched ChatGPT Health. With the usual cheeky self-contradiction intended to avoid legal action, OpenAI invited users to “securely connect medical records and wellness apps to ground conversations in your own health information, so responses are more relevant and useful to you”, while simultaneously warning that ChatGPT Health “is not intended for diagnosis or treatment” (Field, 2026; Muldowney, Hanna, 2026). The *Washington Post* journalist Geoffrey A. Fowler gave the new ChatGPT Health access to a decade of personal data, 29 million steps and 6 million heartbeat measurements stored in his Apple Health app. He then asked the chatbot to grade his cardiac health. The chatbot drew questionable conclusions that changed each time Fowler asked (Fowler, 2026). Of course, the journalist who provided OpenAI with a decade’s worth of heartbeat data in exchange for some magical 8-ball responses was promised the highest standards of health data protection by the company. However, as Nita Farahany and Daniel Solove point out, companies have no fiduciary duty and can change their terms and conditions at any time (Farahany, 2026; Solove, 2026).

The wicked wolf certainly has no intention of curing us, but surely, as he said to Little Red Riding Hood, he is there to listen to us better.

I am grateful to Giuseppe Longo and Maurizio Tirassa for their help and suggestions.

References

- Z. Abrams, *Using generic AI chatbots for mental health support: A dangerous trend*, in *American Psychological Association Services*, March 12, 2025, <https://www.apaservices.org/practice/business/technology/artificial-intelligence-chatbots-therapists>.
- F. Albanese, *A/HRC/59/23: From economy of occupation to economy of genocide - Report of the Special Rapporteur on the situation of human rights in the Palestinian territories occupied since 1967*, June 30, 2025, <https://www.ohchr.org/en/documents/country-reports/ahrc5923-economy-occupation-economy-genocide-report-special-rapporteur>.
- N. Al-Sibai, *Founder of Google’s Generative AI Team Says Don’t Even Bother Getting a Law or Medical Degree, Because AI’s Going to Destroy Both Those Careers Before You Can Even Graduate*, in *Futurism*, August 20, 2025a, <https://futurism.com/former-google-ai-exec-law-medicine>.

- N. Al-Sibai, *ChatGPT Told a Man His Symptoms Were Fine, But Then He Saw a Real Doctor and Realized He Was Dying*, in *Futurism*, September 6, 2025b, <https://futurism.com/chatgpt-symptoms-fine-cancer>.
- S. Altman, A. Huffington, *AI-Driven Behavior Change Could Transform Health Care*, in *Time*, July 7, 2024, <https://time.com/6994739/ai-behavior-change-health-care/>.
- R. Andersen, *Science Is Drowning in AI Slop*, in *The Atlantic*, January 22, 2026, <https://www.theatlantic.com/science/2026/01/ai-slop-science-publishing/685704/>.
- E. Anderson, *Private government. How employers rule our lives (and why we don't talk about it)*, Princeton, Princeton University, 2017.
- Anthropic, *Project Vend: Can Claude run a small shop? (And why does that matter?)*, June 27, 2025, <https://www.anthropic.com/research/project-vend-1>.
- S. Baron-Cohen, *The science of evil. On empathy and the origins of cruelty*, New York, Basic Books, 2011.
- E.M. Bender, T. Gebru, A. McMillan-Major, Sh. Shmitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?* , in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event Canada, ACM, 2021, <https://doi.org/10.1145/3442188.3445922>.
- E.M. Bender, A. Hanna, *AI Con. How to Fight Big Tech's Hype and Create the Future We Want*, London, Penguin Random House, 2025.
- A. Birhane, *Cheap science, real harm: the cost of replacing human participation with synthetic data*, preprint, 2025, <https://synthetic-data-workshop.github.io/papers/13.pdf>.
- A. Blankertz, *Crashing hard: why talking about bubbles obscures the real social cost of overinvesting into "Artificial Intelligence"*, August 12, 2025, <https://www.structural-integrity.eu/crashing-hard-why-talking-about-bubbles-obscures-the-real-social-cost-of-overinvesting-into-artificial-intelligence/>.
- H. Blix, I. Glimmer, *Deflating "Hype" Won't Save Us*, in *LiberalCurrents*, September 16, 2025, <https://www.liberalcurrents.com/deflating-hype-wont-save-us/>.
- P. Bourdieu, *The essence of neoliberalism*, in *Le Monde Diplomatique*, December 1998, <https://mondediplo.com/1998/12/08bourdieu>.
- F. Bria, *Takeover by Big Tech*, in *Le Monde Diplomatique*, November 2025, <https://mondediplo.com/2025/11/02tech>.
- A. Brizio, M. Tirassa, *Biological Agency: Its Subjective Foundations and a Large-Scale Taxonomy*, in *Frontiers in Psychology*, 7, 2016, <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2016.00041>.
- K. Budzyń et alii, *Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study*, in *The Lancet Gastroenterology & Hepatology*, 10, 2025, 896-903, [https://www.thelancet.com/journals/langas/article/PIIS2468-1253\(25\)00133-5](https://www.thelancet.com/journals/langas/article/PIIS2468-1253(25)00133-5).
- C. Cadwalladr, *How to survive the broligarchy: 20 lessons for the post-truth world*, in *The Guardian*, November 17, 2024, <https://www.theguardian.com/commentisfree/2024/nov/17/how-to-survive-the-broligarchy-20-lessons-for-the-post-truth-world-donald-trump>.
- C. Cadwalladr, C. Young, M. Colbert, *Revealed: Palantir deals with UK government amount to at least £670m – including £15m contract with nuclear weapons agency*, in *Foundations of Science*, January 27, 2026, <https://www.thenerve.news/p/palantir-technologies-uk-government-contracts-size-nuclear-deterrent-atomic-peter-thiel-louis-mosley>.
- C.S. Calude, G. Longo, *The Deluge of Spurious Correlations in Big Data*, in *Foundations of Science*, 22, 2017, pp. 595-612, <https://www.di.ens.fr/users/longo/files/BigData-Calude-LongoAug21.pdf>.
- D. Cameron, M. Varner, *ICE and CBP's Face-Recognition App Can't Actually Verify Who People Are*, in *Wired*, February 5, 2026, <https://www.wired.com/story/cbp-ice-dhs-mobile-fortify-face-recognition-verify-identity/>.
- K. Caridad, *When the Chatbot Becomes the Crisis: Understanding AI-Induced Psychosis*, July 1, 2025, <https://www.papsychotherapy.org/blog/when-the-chatbot-becomes-the-crisis-understanding-ai-induced-psychosis>.

R. Caso, P. Guarda, *Ricerca e Spazio Europeo dei Dati Sanitari tra regole proprietarie e apertura*, in *Accademia - Rivista dell'Associazione dei Civilisti Italiani*, 9, 2025, 837-850, <https://zenodo.org/records/18253933>.

A. Challapally, C. Pease, R. Raskar, P. Chari, *The GenAI Divide. State of AI in Business 2025*, MIT Nanda, July 2025, https://mlq.ai/media/quarterly_decks/vo.1_State_of_AI_in_Business_2025_Report.pdf.

D. Coldewey, *Signal's Meredith Whittaker: AI is fundamentally 'a surveillance technology'*, in *TechCrunch*, September 25, 2023, <https://techcrunch.com/2023/09/25/signals-meredith-whittaker-ai-is-fundamentally-a-surveillance-technology/>.

A. Coveri, C. Cozza, D. Guarascio, *Big Tech and the US Digital-Military-Industrial Complex*, in *Intereconomics*, 60, 2, 2025, 81-87, <https://sciendo.com/article/10.2478/ie-2025-0017>.

J. Cox, *Here is the Agreement Giving ICE Medicaid Patients' Data*, in *404 Media*, January 6, 2026a, <https://www.404media.co/here-is-the-agreement-giving-ice-medicaid-patients-data/>.

J. Cox, *'ELITE': The Palantir App ICE Uses to Find Neighborhoods to Raid*, in *404 Media*, January 15, 2026b, <https://www.404media.co/elite-the-palantir-app-ice-uses-to-find-neighborhoods-to-raid/>.

K. Crawford, *Atlas of AI. Power, Politics, and the Planetary Costs of Artificial Intelligence*, New Haven and London, Yale University Press, 2021.

J. De Freitas, Z. Oguz-Uguralp, A. Kaan-Uguralp, *Emotional Manipulation by AI Companions*, last revised October 7, 2025, <https://arxiv.org/abs/2508.19258>.

J.C. De Martin, *The computerization of the world and international cooperation*, in *Nexa Center for Internet & Society*, December 2024, <https://nexa.polito.it/the-computerization-of-the-world-and-international-cooperation/>.

C. Doctorow, *The Shitty Tech Adoption Curve Has a Business Model*, in *Pluralistic*, June 11, 2023, <https://pluralistic.net/2023/06/11/the-shitty-tech-adoption-curve-has-a-business-model/>.

C. Doctorow, *The real (economic) AI apocalypse is nigh*, in *Pluralistic*, September 27, 2025, <https://pluralistic.net/2025/09/27/econopocalypse/#subprime-intelligence>.

A.C. Doyle, *The Adventure of Silver Blaze*, in *The Strand Magazine*, 24, 1892, 645-660, <https://archive.org/details/StrandMagazine24/page/n95>.

A. Eichenberger, S. Thielke, A. Van Buskirk, *A Case of Bromism Influenced by Use of Artificial Intelligence*, in *Annals of Internal Medicine: Clinical Cases*, 4, 8, 2025, <https://doi.org/10.7326/aimcc.2024.1260>.

N. Farahany, *Your Doctor Has a Fiduciary Duty to You. ChatGPT Doesn't*, January 14, 2026, <https://staytuned.substack.com/cp/184600139>.

G. Farruggia, *AI can transform healthcare – if we transform our data architecture*, in *World Economic Forum*, <https://www.weforum.org/stories/2026/01/ai-healthcare-data-architecture/>.

L. Feiger, S. Levy, *Dr. Oz Pushed for AI Health Care in First Medicare Agency Town Hall*, in *Wired*, April 8, 2025, <https://www.wired.com/story/dr-oz-ai-health-care-medicare-cms-town-hall/>.

H. Field, *OpenAI launches ChatGPT Health, encouraging users to connect their medical records*, in *The Verge*, January 7, 2026, <https://www.theverge.com/ai-artificial-intelligence/857640/openai-launches-chatgpt-health-connect-medical-records>.

R. Fellner, E. Sutherland, K. Vujovic, *Digitalisation of public health: Leading practices in immunisation reporting and respiratory disease surveillance*, OECD Health Working Papers, 185, Paris, OECD Publishing, 2025, <https://doi.org/10.1787/97204768-en>.

F. Floricic, *Non-dits et dommages collatéraux de l'IA générative*, in *Texto! Textes et cultures*, 30, 3-4, 2025, https://www.revue-texto.net/docannexe/file/511/texto_floricic_ia_non_dits_dommages.pdf.

E. Forlini, *Vibe Coding Fiasco: AI Agent Goes Rogue, Deletes Company's Entire Database*, in *PCMag*, July 22, 2025, <https://www.pcmag.com/news/vibe-coding-fiasco-replite-ai-agent-goes-rogue-deletes-company-database>.

G.A. Fowler, *I let ChatGPT analyze a decade of my Apple Watch data. Then I called my doctor*, in *The Washington Post*, January 26, 2026, <https://www.washingtonpost.com/technology/2026/01/26/chatgpt-health-apple/>.

C. Ginzburg, *Clues: Roots of an Evidential Paradigm*, in Idem, *Clues, myths, and the historical method*, translated by J. Tedeschi, A.C. Tedeschi, Baltimore, John Hopkins University Press, 1989, 96-125, <https://archive.org/details/cluesmythshistoroogoginz/page/96>.

R.J. González, *How Big Tech and Silicon Valley are Transforming the Military-Industrial Complex*, in *Thomas J. Watson Jr. School of International and Public Affairs*, Brown University, April 17, 2024, <https://watson.brown.edu/costsofwar/papers/2024/SiliconValley>.

L. Goodridge, O. Blackstock, *We must not let AI 'pull the doctor out of the visit' for low-income patients*, in *The Guardian*, January 25, 2026, <https://www.theguardian.com/commentisfree/2026/jan/25/ai-healthcare-risks-low-income-people>.

J. Green, *When will the AI bubble burst?*, in *TechHQ*, March 10, 2025, <https://techhq.com/news/will-the-ai-bubble-burst-when-will-artificial-intelligence-market-crash/>.

A. Gregory, *Google AI Overviews put people at risk of harm with misleading health advice*, in *The Guardian*, January 2, 2026, <https://www.theguardian.com/technology/2026/jan/02/google-ai-overviews-risk-harm-misleading-health-information>.

O. Guest, A.E. Martin, *A Metatheory of Classical and Modern Connectionism*, in *Psychological Review*, 2025a, https://osf.io/preprints/psyarxiv/eaf2z_v1.

O. Guest, A.E. Martin, *Are Neurocognitive Representations 'Small Cakes'?*, in *PhilSci-Archive*, 2025b, <https://philsci-archive.pitt.edu/24834>.

O. Guest, I. van Rooij, *Critical Artificial Intelligence Literacy for Psychologists*, preprint, October 3, 2025c, https://osf.io/preprints/psyarxiv/dkrgj_v1.

J. Hendrix, *Breaking Down the Lawsuit Against OpenAI Over Teen's Suicide*, in *Tech policy Press*, August 27, 2025, <https://www.techpolicy.press/breaking-down-the-lawsuit-against-openai-over-teens-suicide/>.

G. Hinton, *On Radiology*, 2016, <https://yewtu.be/watch?v=2HMPRXstSvQ>.

G. Huckins, *This medical startup uses LLMs to run appointments and make diagnoses*, in *MIT Technology Review*, September 22, 2025, <https://web.archive.org/web/20250922125538/https://www.technologyreview.com/2025/09/22/1123873/medical-diagnosis-llm/>.

M. Jones, *The CDC, Palantir and the AI-Healthcare Revolution*, in *Unlimited Hangout*, January 13, 2025, <https://unlimitedhangout.com/2025/01/investigative-reports/the-cdc-palantir-and-the-ai-healthcare-revolution/>.

I. Kant, *Vorarbeiten zu Zum Ewigen Frieden (VZeF)*, in *Kant's Gesammelte Schriften*, Prussian Academy Edition, Berlin, Walter de Gruyter and predecessors, 1900-, <https://www.korpora.org/kant/aa23/175.html>.

A. Kaplan, *The Conduct of Inquiry. Methodology for Behavioral Science*, San Francisco, Chandler Publishing Company, 1964, <https://archive.org/details/conductofinquiryoooookapl/page/28/mode/2up>.

M. Kaufman, *'No shortage of wealth': Mamdani, other Democrats chide NYC hospital executives over nurses' strike*, in *Politico*, January 13, 2026, <https://www.politico.com/news/2026/01/13/no-shortage-of-wealth-lawmakers-chide-nyc-hospital-executives-over-nurses-strike-00724577>.

T. Kim, *Goldman Sachs asks in biotech research report: 'Is curing patients a sustainable business model?'*, in *CNBC*, April 11, 2018, <https://www.cnbc.com/2018/04/11/goldman-asks-is-curing-patients-a-sustainable-business-model.html>.

R. Kitchin, *The Data Revolution. Big Data, Open Data, Data Infrastructures and Their Consequences*, Los Angeles, Sage Publications, 2014.

N. Lambert, L. Castricato, L. von Werra, A. Havrilla, *Illustrating Reinforcement Learning from Human Feedback (RLHF)*, December 9, 2022, <https://huggingface.co/blog/rlhf>.

F. Landymore, *OpenAI CEO Says AI Will Give Medical Advice to People Too Poor to Afford Doctors*, in *Futurism*, February 21, 2023, <https://futurism.com/the-byte/openai-ceo-ai-medical-advice>.

G. Longo, *The reasonable effectiveness of Mathematics and its Cognitive roots*, in *Geometries of Nature, Living Systems and Human Cognition series in New Interactions of Mathematics with Natural Sciences and Humanities*, 2005, 351 - 382, <https://www.di.ens.fr/users/longo/files/PhilosophyAndCognition/reason-effect.pdf>.

G. Longo, M. Montévil, *From logic to biology via physics: a survey*, "Logical Methods in Computer Science", 13, 4, 2017, <https://hal.science/hal-01377051v2/document>.

G. Longo, *Information at the Threshold of Interpretation Science as Human Construction of Sense*, in *A Critical Reflection on Automated Science – Will Science Remain Human?*, ed. by M. Bertolaso, F. Sterpetti, Springer, 2020, 67-99, https://hal.science/hal-02903688v1/file/_2019_Information-Interpretation.pdf.

G. Longo, *Le cauchemar de Prométhée. Les sciences et leurs limites*, Paris, Presses Universitaires de France, 2023a.

G. Longo, *Emergence vs Novelty Production in Physics vs Biology*, Lecture presented at Open Historicity of Life. Theory, epistemology, practice, Paris, October, 2023b, to appear in the proceedings, ed. by M. Chollat-Namy, M. Montévil, A. Robert, <https://www.di.ens.fr/users/longo/files/EmergeCompareBioNovelty.pdf>.

G. Longo, *La créativité du vivant face à l'émergence des corrélations: épistémologie et politique des Large Language Models*, in *Penser l'intelligence artificielle. Enjeux philosophiques, culturels et politiques de l'automatisation numérique*, ed. by Anne Alombert, Alban Leveau-Vallier, Baptiste Loreaux., Paris, Presses du Réel, 2025, <https://www.di.ens.fr/users/longo/files/LLM-emergence-creativite.pdf>

J. MacDonald-Evoy, *ICE director envisions Amazon-like mass deportation system: 'Prime, but with human beings'*, in *Arizona Mirror*, April 8, 2025, <https://azmirror.com/2025/04/08/ice-director-envisions-amazon-like-mass-deportation-system-prime-but-with-human-beings/>.

A. MacLeod, *Meet the ex-CIA agents deciding Facebook's content policy*, in *MintPress News*, July 12, 2025, <https://www.mintpressnews.com/meet-ex-cia-agents-deciding-facebook-content-policy/281307/>.

D. Mcquillan, *AI as Algorithmic Thatcherism*, December 21, 2023, https://www.danmcquillan.org/ai_thatcherism.html.

E. Maiberg, *AI-Generated Science*, in *404 Media*, March 18, 2024, <https://www.404media.co/email/a2a944f8-235a-4c75-8d00-955edbbfcb4e/>.

S. Maurizi, January 21, 2026, <https://poliversity.it/@smaurizi@mastodon.social/115933342753230715>

L. Messeri, M.J. Crockett, *Artificial intelligence and illusions of understanding in scientific research*, in *Nature*, 627, 2024, 49-58, <https://lrc.northwestern.edu/language-instruction/professional-development/artificial-intelligence-and-illusions-of-understanding-in-scientific-research-nature.pdf>.

Meta, *Improving Your Recommendations on Our Apps With AI at Meta*, October 1, 2025, <https://about.fb.com/news/2025/10/improving-your-recommendations-apps-ai-meta/>.

M.R. Mhawish, *Watched, Tracked, and Targeted. Life in Gaza under Israel's all-encompassing surveillance regime*, in *Intelligencer*, December 3, 2025, <https://nymag.com/intelligencer/article/watched-tracked-targeted-israel-surveillance-gaza.html>.

E. Morrow, S. Hopewell, E. Williamson, T. Theologis, *Threat of imposter participants in health*, in *Bmj*, 391, 2005, <https://www.bmj.com/content/391/bmj.r2128>.

D. Mousa, *The algorithm will see you now*, in *Works in Progress*, September 25, 2025, <https://worksinprogress.co/issue/the-algorithm-will-see-you-now/>.

D. Muldowney, A. Hanna, *ChatGPT Wants Your Health Data*, January 16, 2026, <https://buttondown.com/maiht3k/archive/chatgpt-wants-your-health-data/>.

- H. Nolan, *Gangster Party*, in *How Things Work*, February 21, 2025a, <https://www.hamiltonnolan.com/p/gangster-party>.
- H. Nolan, *Welcome to the Age of Disappearance*, in *How Things Work*, July 2, 2025b, <https://www.hamiltonnolan.com/p/welcome-to-the-age-of-disappearance>.
- T. Numerico, *Big data e algoritmi. Prospettive critiche*, Roma, Carocci, 2021.
- M. O'Brien, *Musk's latest Grok chatbot searches for billionaire mogul's views before answering questions*, *The Associated Press*, July 11, 2025, <https://apnews.com/article/grok-4-elon-musk-xai-colossus-14d575fb490c2b679ed311a1c83f857>.
- J.D. Oliva, *Regulating Healthcare Coverage Algorithm*, in *Indiana Law Journal*, 100, 4, 2025, <https://www.repository.law.indiana.edu/cgi/viewcontent.cgi?article=11586&context=ilj>.
- OpenAI, *GPT-4 Technical Report*, March 15, 2023, <https://arxiv.org/abs/2303.08774v1>.
- M. Padden, *The transformation of surveillance in the digitalisation discourse of the OECD: a brief genealogy*, in *Internet Policy Review*, 12, 3, 2023, <https://doi.org/10.14763/2023.3.1720>.
- A. Parshall, *Why ChatGPT Shouldn't Be Your Therapist*, in *Scientific American*, August 13, 2025, <https://www.scientificamerican.com/article/why-ai-therapy-can-be-so-dangerous/>.
- S. Perez, *Signal President Meredith Whittaker calls out agentic AI as having 'profound' security and privacy issues*, in *Tech Crunch*, March 7, 2025, <https://techcrunch.com/2025/03/07/signal-president-meredith-whittaker-calls-out-agentic-ai-as-having-profound-security-and-privacy-issues/>.
- N. Perret, G. Longo, *Reductionist perspectives and the notion of information*, in *Progress in Biophysics and Molecular Biology*, 122, 1, 2016, 11-15, https://www.di.ens.fr/users/longo/files/o2_information.pdf.
- R. Picardi, *Cosmopolitan right, faraway injustices and peace*, forthcoming in *War and Peace after Kant*, ed. by M. Hergouth, L. Kuhar, Bloomsbury Press, 2026.
- M.C. Pievatolo, *Improprietà intellettuale. Ricostruire il diritto d'autore*, in *Bollettino telematico di filosofia politica*, 2024, <https://btfp.sp.unipi.it/dida/cristallo/index.xhtml>.
- C. Rapezzi, R. Ferrari, A. Branzi, *White coats and fingerprints: diagnostic reasoning in medicine and investigative methods of fictional detectives*, in *Bmj*, 331, 2005, 1491-1494, <https://www.bmj.com/content/331/7531/1491>.
- A. Rapp, M. Tirassa, *Know Thyself: A Theory of the Self for Personal Informatics*, in *Human-Computer Interaction*, 32, 5-6, 2017, 335-380, <https://courses.grainger.illinois.edu/cs565/sp2018/Lives.pdf>.
- E. Reinhart, *What we lose when we surrender care to algorithms*, in *The Guardian*, November 9, 2025, <https://www.theguardian.com/us-news/ng-interactive/2025/nov/09/healthcare-artificial-intelligence-ai>.
- M. Ricciardi, *È per noi la pubblicità dei droni killer*, in *Il Manifesto*, 17 luglio 2025, <https://ilmanifesto.it/e-per-noi-la-pubblicita-dei-droni-killer>.
- J. Rigby, *Gates and OpenAI team up for AI health push in African countries*, January 21, 2026, in *Reuters*, <https://www.reuters.com/business/healthcare-pharmaceuticals/gates-openai-team-up-ai-health-push-african-countries-2026-01-21/>.
- E. Rothermich, *Hospice Commodification and the Limits of Antitrust*, in *Law and Political Economy (LPE) Project*, November 12, 2025, <https://lpeproject.org/blog/hospice-commodification-and-the-limits-of-antitrust/>.
- P. Rucker, *How Cigna Saves Millions by Having Its Doctors Reject Claims Without Reading Them*, in *ProPublica*, March 25, 2023, <https://www.propublica.org/article/cigna-pdx-medical-health-insurance-rejection-claims>.
- P. Rumale V et alii, *Faithfulness Hallucination Detection in Healthcare AI*, in *Proceedings of KDD 2024 Workshop - Artificial Intelligence and Data Science for Healthcare: Bridging Data-Centric AI and People-Centric Healthcare*, New York, ACM, preprint, 2025, <https://openreview.net/pdf?id=6eMIzKFOPJ>.
- J. Sadowski, *The Mechanic and the Luddite. A Ruthless Criticism of Technology and Capitalism*, Oakland, California, University of California Press, 2025.
- E. Salvaggio, *Human conversation*, September 7, 2025, <https://mail.cyberneticforests.com/human-conversation/>.

- E. Saner, 'Nobody wants a robot to read them a story!' *The creatives and academics rejecting AI – at work and at home*, in *The Guardian*, June 3, 2025, <https://www.theguardian.com/technology/2025/jun/03/creatives-academics-rejecting-ai-at-home-work>.
- S. Schechner, S. Kessler, 'I Feel Like I'm Going Crazy': ChatGPT Fuels Delusional Spirals, in *The Wall Street Journal*, August 7, 2025, <https://www.wsj.com/tech/ai/i-feel-like-im-going-crazy-chatgpt-fuels-delusional-spirals-aes5rfc>.
- B. Schneier, *What LLMs Know About Their Users*, June 25, 2025, <https://www.schneier.com/blog/archives/2025/06/what-llms-know-about-their-users.html>.
- J. Shah, N. Shafi, C. Rallo, *Announcing agentic AI for healthcare patient engagement in Amazon Connect*, November 26, 2025, <https://aws.amazon.com/it/blogs/industries/announcing-agentic-ai-for-healthcare-patient-engagement-in-amazon-connect-preview/>.
- M. Shanahan, *Talking about Large Language Models*, in *Communications of the ACM*, 67, 2, 2024, 68–79, <https://doi.org/10.1145/3624724>.
- J. Singley, "We Couldn't Generate an Answer for your Question", in *ACRLog*, July 21, 2025, <https://acrlog.org/2025/07/21/we-couldnt-generate-an-answer-for-your-question/>.
- D. Solove, *Privacy in Authoritarian Times: Surveillance Capitalism and Government Surveillance*, forthcoming in *Boston College Law*, 2025, <https://www.almendron.com/tribuna/wp-content/uploads/2025/05/ssrn-5103271.pdf>.
- D. Solove, *AI Companies Should Have Information Fiduciary Duties*, January 15, 2026, <https://danielsolove.substack.com/p/ai-companies-should-have-information>.
- J. Sooknanan, T. Seemungal, *Not so elementary - the reasoning behind a medical diagnosis*, in *MedEdPublish*, 8, 17 December 2019, <https://pubmed.ncbi.nlm.nih.gov/38089286/>.
- J. Stokes, *ChatGPT explained: a normie's guide to how it works*, March 1, 2023, <https://www.jonstokes.com/p/chatgpt-explained-a-guide-for-normies>.
- V. Tangermann, *Therapy Chatbot Tells Recovering Addict to Have a Little Meth as a Treat*, in *Futurism*, June 2, 2025a, <https://futurism.com/therapy-chatbot-addict-meth>.
- V. Tangermann, *Doctors Horrified After Google's Healthcare AI Makes Up a Body Part That Does Not Exist in Humans*, in *Futurism*, August 6, 2025b, <https://futurism.com/neoscope/google-healthcare-ai-makes-up-body-part>.
- M. Tirassa, *Intelligenza artificiale e mondi reali*, in *Nexa Center for Internet & Society*, marzo 2025, <https://nexa.polito.it/intelligenza-artificiale-e-mondi-reali/>.
- I. van Rooij et alii, *Reclaiming AI as a Theoretical Tool for Cognitive Science*, «Computational Brain & Behavior», 2024, <https://doi.org/10.1007/s42113-024-00217-5>.
- I. van Rooij, *AI slop and the destruction of knowledge*, August 12, 2025, <https://irisvanrooijcogsci.com/2025/08/12/ai-slop-and-the-destruction-of-knowledge/>.
- J. Au Yeung et alii, *The Psychogenic Machine: Simulating AI Psychosis, Delusion Reinforcement and Harm Enablement in Large Language Models*, September 13, 2025, arxiv.org/pdf/2509.10970.
- S.M. Yoon et alii, *Navigating artificial intelligence (AI) accuracy: A meta-analysis of hallucination incidence in large language model (LLM) responses to oncology questions*, in *Journal of Clinical Oncology*, 43, June 2025, https://ascopubs.org/doi/pdf/10.1200/JCO.2025.43.16_suppl.e13686.
- C. Warzel, *AI Has Become a Technology of Faith*, in *The Atlantic*, July 12, 2024, <https://www.theatlantic.com/technology/archive/2024/07/thrive-ai-health-huffington-altman-faith/678984>.
- C. Warzel, *AI Is a Mass-Delusion Event*, in *The Atlantic*, August 18, 2025, <https://www.theatlantic.com/technology/archive/2025/08/ai-mass-delusion-event/683909/>.

E. Webb, S. Goel, *Leaked docs show how Meta is training its chatbots to message you first, remember your chats, and keep you talking*, in *Business Insider*, July 3, 2025, <https://www.businessinsider.com/meta-ai-studio-chatbot-training-proactive-leaked-documents-alignerr-2025-7>.

S. Wells, *Exploring the Dangers of AI in Mental Health Care*, in *Stanford HAI*, June 11, 2025, <https://hai.stanford.edu/news/exploring-the-dangers-of-ai-in-mental-health-care>.

J. Wilkins, *OpenAI Reportedly Planning to Make ChatGPT “Prioritize” Advertisers in Conversation*, in *Futurism*, December 30, 2025, <https://futurism.com/artificial-intelligence/openai-chatgpt-sponsored-ads>.

D. Wu, *Illinois bans AI therapy as some states begin to scrutinize chatbots*, in *The Washington Post*, August 12, 2025, <https://www.washingtonpost.com/nation/2025/08/12/illinois-ai-therapy-ban/>.

A.L. Zimmerman, *Evidence-Based Medicine: A Short History of a Modern Medical Movement*, in *Virtual Mentor. American Medical Association Journal of Ethics*, 15, 1, January 2013, 71-76, <https://journalofethics.ama-assn.org/sites/joedb/files/2018-06/mhstr-1301.pdf>.

