An Ethics and Security Challenge for Oncology: Communicate the Implementation of Generative AI

https://doi.org/10.32635/2176-9745.RBC.2025v71n2.5184EN

Um Desafio de Ética e de Segurança para a Oncologia: Comunicar a Implementação da IA Generativa

Un Desafío de Ética y de Seguridad para la Oncología: Comunicar la Implementación de la IA Generativa

Nemézio Clímico Amaral Filho[1]

INTRODUCTION

Most companies are not prepared to deal with cyberattacks. This statement has gained new implications in 2023 with the use of generative artificial intelligence (AI). In Brazil that year, 54% of companies targeted by cyberattacks had generative AI to blame[1]. An investigation encompassing 170 countries shows the situation has worsened in 2024: information theft using infostealer malware has grown 58%. Moreover, the authors claimed the greatest victims of cyberattacks were education (75%), health (47%), and government (43%) spheres[1]. The health industry was the second most targeted by ransomware (software that holds victims' confidential data or devices hostage, keeping them blocked or threatened by destruction if a ransom is not paid) manipulators[2].

The popularity and absence of global regulation on generative and non-generative AI has fostered the creativity of cybercriminals. Cyberattacks are serious but scarcely the only problems faced by public and private healthcare services brought on by unregulated and uncontrolled development of AI, which is being applied to treatment and research, whether society, the latest target of such innovation, is aware of it or not. How can we ensure the ethical efficacy of a technology application that uses human health as an empirical field, allowing active transparency and communication of the process to be verified and perceived by society?

The issue is part of a debate that is slowly gaining ground with different approaches in communication media in the shape of a question that precedes the one formulated in the previous paragraph and is increasingly unavoidable: was society prepared to absorb this innovation with no prior security safeguards? Apparently, given the damage done to democracy across the West, most citizens were not even ready to use a smartphone as a "freedom of expression" instrument without previous training or mediation. This leads to other issues in the case of AI in the healthcare system: how can we ensure patients' and research data are not stolen or adulterated? Also, what is the state of security concerning the application of such technology in this field? Is it possible that health professionals may over-rely on the machine? These questions forcibly lead to more incisive questioning in specific fields of medicine, such as cancer research. For instance, how does oncology intend to communicate the issue to society in a straightforward, responsible manner to reduce the risks of using such a deregulated technology?

Until the drafting of this article, no one seemed to have objective answers to those questions, which is perfectly reasonable in the legally and anarchically political environment in which AI in general, and generative AI, particularly, have been thriving. But it is possible to predict some risks to support the argument that public healthcare must understand the complex nature of the problem and, more importantly, understand the technical, political, and social viewpoints, as their representatives will soon be required to communicate with society. In both intelligible and lay terms.

It is important to take a few steps back before committing to the ethical use of generative AI in oncology and the need for transparent communication of this practice to society.

DEVELOPMENT

To illustrate our point, in early 2025, Google developed an AI tool that works as a virtual collaborator of biomedical scientists. This new tool, tested by scientists of Stanford University (USA) and Imperial College London[3] (England), uses advanced "reasoning" to help scientists summarize large amounts of scientific texts and generate

[1]Instituto Nacional de Câncer (INCA). Universidade do Estado do Rio de Janeiro (Uerj). Rio de Janeiro (RJ), Brasil. Orcid: https://orcid.org/0000-0002-3902-2800
Corresponding author: Nemézio Clímico Amaral Filho. Rua Paissandu, 156/604 – Flamengo. Rio de Janeiro (RJ), Brasil. CEP 22210-080. E-mail: nfilho@inca.gov.br

new hypotheses. In an experiment on hepatic fibrosis, Google ensured that *all* the approaches suggested by its new AI "co-scientist" (as the company referred to their model) showed promising activity and potential to inhibit disease causes. Google "expects" its tool to aid and not decrease scientific collaboration.

According to the Brazilian Federal Council of Biomedicine[4]:

> The Biomedical professional is responsible for identifying, classifying, and studying microorganisms that cause infirmities. Develop medications and produce vaccines to counter diseases. Test and interpret the results of clinical analysis to diagnose diseases, verify contamination in food.

Too much responsibility for a machine, skeptics could argue, which is why defining its application could only be performed by a human being. But without effective supervision, active communication of tools, capabilities, and decisive limits of generative AI software applied to the medical field, when pressured by time and results, would some professionals not be tempted to delegate sensitive analyses to machines?

The promise of an AI tool that can make decisions and/or strong suggestions, working as an auxiliary, prevails in technology companies' discourse. In this narrative, qualified professionals would thus be liberated to work in the most important aspects of their roles. From a marketing standpoint, it would be suicide not to promote this idea. Otherwise, AI companies would be promising a technology that replaces centuries of accumulated human knowledge and experience in the medical field, going against a strongly rooted belief in the collective world mind. Despite believing in the potential of their technology to replace many jobs occupied by once sacred professionals, companies know they would face innumerable potential lawsuits, since errors are inevitable – error is actually what drives development, knowledge. And if errors must exist, let them be admitted by humans and not their assistants or "co-scientists", in this example of AI applied to biomedicine.

There are many other shades of gray missing in the co-assistance and AI adoption discourse in several sensitive fields in people's daily lives. When a software that calculated the probability of an offender committing crimes again in the USA, therefore helping the judge decide if they should be granted probation or not, was adopted in the legal system, it seemed like a good option. It was an "unbiased" machine calculating who deserved a second chance. Thus, the Correctional Offender Management Profiling for Alternative Sanctions (Compas) was adopted in some USA States. In 2016, investigative journalists from the ProPublica[5] organization uncovered that the Compas rate of correct predictions in 2013-2014 was 60%, which means that society or recoverable suspects were undoubtedly harmed in 40% of predictions simply because judges delegated their decision to the algorithm, despite it being only an "assistant" of the Justice system. Black people received the worst ratings because the algorithm was biased against this population, like many are. The company that produced it, Northpointe Inc., did not disclose its rating system.

Let there be no doubts: both the biomedicine domain and the Justice examples deal with social trust, transparency, and security. Or the lack thereof, with a rather cynical excuse that adjustments are made based on "user experience". Among so many coincidences, these software differ only in their algorithmic structure. Compas is a "hierarchical decision" or "rule-based" model, whereas Google's co-scientists is a Transformer application, from which derive the most recent generative AI models from GPT to DeepSeek. The objective, however, is the same – to convince of its effectiveness.

Google's co-scientist surprised again part of the audience more tuned to the scientific promises of Large Language Models (LLM): Imperial College microbiologists spent a decade drafting, testing, and proving hypotheses – which, in the world of laws, norms, and ethics, by itself requires time – to understand why some super bacteria are immune to antibiotics. Google's co-scientist received a brief question about what microbiologists researched, and the machine reached the same conclusion as humans in 48 hours. A mixed feeling hit researchers: on one hand, they asked themselves if they remained relevant; on the other hand, how could they use the powerful tool in their hands to unlock even more advanced research[6].

Considering all these models, the ethical question to be asked is: can we, should we bestow so much power to deregulated, developing algorithmic structures surrounded by industrial secrets? We can mention the effects in any field of human knowledge, but it causes even more heightened ethical tensions in the health field.

The argument that medical professionals cannot be replaced by machines is frequent. And this certainty, to its advocates, is all the communication they really need – everything else is a matter of adjustment. Truth be told, companies at the forefront of generative AI, still captained by the USA – but with China coming in strong, strengthening the perception that overcoming their competitor is a matter of (little) time – echo this discourse due to the reasons described, among others. Why these marketing strategies – of which the AI concept itself originates from, and not necessarily a scientific

field – interest big technology corporations was further addressed elsewhere[7].

It is important to highlight that not all AI researchers are as subtle in their commercial discourse on the application of emerging technology in health. China has been developing an algorithm like the one advertised by Google for biomedicine for ten years before the American giant. The RXThinking startup trains medical AI algorithms to become "super diagnosers" [8], and that should be further enhanced with the union of generative AI with the so-called "autonomous intelligent agents", who can make decisions for the users, from planning a trip and buying plane tickets to informing a lack of medication.

Google[9] published a report predicting the expansion of AI agents in different areas by 2025, including healthcare. The "multimodal AI" that imitates human learning by cross-referencing several data sources would also become more frequent in complex data analyses in the health field. To Google, adopting multimodal AI models comes closer to the idea of "personalized medicine". The technology would also serve to reduce the administrative work performed by doctors.

However, back in 2019, Lee, former president of Google China, Microsoft, and Apple executive, now one of the people thinking about the next generation of innovation companies in China, bravely put the cards on the table:

> Despite the many social elements a doctor's appointment represents, the diagnostic core involves data collection (symptoms, medical history, environmental factors) and forecasting correlated phenomena (a disease). Searching for several correlations and making forecasts is exactly what deep learning does best. With sufficient training data – in that case, precise medical records –, an AI-powered diagnosis tool may transform any medical professional into a superdiagnoser, a doctor with the experience of tens of millions of cases, an unusual ability to detect concealed correlations, and a perfect memory to be leveraged[8].

Kai-Fu[8] has more than machine learning predictions to sustain his prognosis. In 2013, he was diagnosed with cancer, a lymphoma defined as stage IV by the first doctor who assisted him, after a PET scan analysis that allowed this as the only interpretation (this image test assesses the function and presence of diseases in organs and tissues). His lawyer even advised him to draft his will, but the scientist found out the metrics used were not enough to determine his cancer stage and decided to search on his own for studies that considered other variables.

He found the work of Italian researchers that questioned the variables used in the PET scan, finding markers closely related to his case. The initial diagnosis represented a 50% five-year survival rate, but the more detailed item of the Italian scientific article increases this percentage to 89%. A new doctor, the greatest lymphoma specialist in Taiwan, confirmed that designating Kai-Fu's cancer as stage IV was an error. The scientist's cancer went into remission, and until the draft of this text, he was still alive and researching for the next AI wave[8].

> The reality is that it will not take long for AI algorithms to perform many of the diagnostic functions performed by medical professionals. These algorithms will identify diseases and prescribe treatments more effectively than any human being. In some cases, doctors will use these equations as a tool. In other cases, the algorithms may completely replace doctors.

In 2019, a study[10] published in The Lancet Digital Health, conducted by researchers of the University of Birmingham, evaluated diagnoses obtained from AI image analysis since 2012, noting they were comparable to the ones by clinical specialists – reinforcing Kai-Fu Lee's prediction.

The Chinese researcher considers, however, that an application does not replace a doctor, who, if they wish, can ignore millions of health records searched by software, as well as the latest medical publications, finally basing the diagnosis on their own decisions. But, liberated by technology, doctors could do things machines cannot. For instance, comfort patients and families when a diagnosis is unfavorable. However, in the era of generative AI, even affection can be "trained" by algorithms.

In 2023, a study[11] published in JAMA Internal Medicine showed that ChatGPT potentially gave extremely acceptable answers to routine medical questions. Moreover, 78% of the assessed informed that the machine's answers were more empathetic than the doctors'. The chatbot would be delivering more "affection" and additional care with words. This could happen due to professionals' tribulation, tiredness, or sheer lack of tact.

In interactions with the middle class, skepticism and rejection are noted, mainly regarding the cordial care the patient may receive with the comforting word of a chatbot. In the user-patient perception, this more affectionate treatment is not always dispensed by doctors. It is so easy to find someone in Brazil nowadays who has cancer, or a relative or friend with the disease, as it is to find someone who was treated with discomfort by a health professional at least once in their lives, which could partly explain the good impression transmitted by the chatbot.

This popular perception and scientific studies converge on the same spot: social and professional reality was altered by AI technology, and the sacred aspect of medical activity is being questioned. And that, for better or worse, is also a question of security that public and private health systems must face pragmatically to elaborate a communication strategy for a society disputed by competitor discourses whose weights are not exclusively attributed by science.

However, what this article defends is that generative AI applied to several areas of knowledge, especially healthcare – from management to clinical analysis –, cannot be faced in a luddite manner. Recently, researchers of the Brazilian National Cancer Institute (INCA)[12] and other health institutions in Brazil defended the use of tools such as ChatGPT to democratize knowledge for people with cancer. None of them ignore that these AIs can be imprecise and provide harmful information, which increases the importance of AI education, helping demystify the technology and, thus, obtain safer information. In contrast, the enormous capability of generative AI to analyze and compare data enables oncological studies that would be impossible until recently.

The study "Exploring the role of ChatGPT in patient care (diagnosis and treatment) and medical research: a systematic review"[13] reviewed over a hundred articles to assess the use of this generative AI in oncology and concluded: "Although it can help with patient treatment and research, there are issues with accuracy, authorship, and bias. ChatGPT can serve as a 'clinical assistant' and help in research and scholarly writing." The article also shows that AI can clarify information for less favored individuals without being offensive. But there is a risk of oversimplification and underestimation of the severity of a disease.

And, as one can figure, risks are not limited to errors due to biased machine data or hardware limitations. What humans consciously decide when introducing the technology as a "business model" matters a lot, which, in the healthcare case, may not even mean direct applications in the field. Justpoint[14] startup is a good example. Based in Colorado, USA, it holds expertise in AI, pharmaceutical products, and the legal system. It is a self-declared consumer protection company. Between financing and credit, Justpoint raised 95 million dollars to build a structure that uses AI and scientists to search medical records to find medication, chemical, and utility products that, according to it, have harmed patients and consumers. This information is then provided to a law firm in the neighboring state of Arizona, Justpoint PLLC, that can prosecute the manufacturers of those products and have Justpoint Inc. keep a share of the successful earnings. In the name of consumers and patients, the organization intends to increase earnings in the billion-dollar market of American lawsuit actions against the equally billionaire pharmaceutical industry.

As reasonable people, including health professionals, are well aware, there is a fine line between what is medication and what is poison, and what is good for one body may not be good for another. How to undo the possible damage caused to the prosecuted brand that acted in good faith, the doctor who prescribed it, and any other patient who may have received the medication, be denied it due to fear of an alleged error being replicated? Should this really be in the hands of a startup "business model", as well-intentioned as it may be, or should this be part of a public, enlightening debate to provide the relevant legislator with reliable information?

Some organizations worldwide are already helping build arguments for the necessary legal safeguards that will have to be put into place by the legislative branch of countries mindful of the issue. In the medical sphere, an example of such an organization is the Australian Medical Association (Ama)[15], which has published its stance on the use of AI in healthcare. Among other statements, Ama advocates that AI supports doctors and health professionals in general in meeting the community's needs: the technology may improve but not replace clinical decision-making. In this interpretation, AI should never compromise doctors' clinical independence and professional autonomy. The organization sustains that employing AI in healthcare must not increase disparities in any population segment. It also advocates that the use of AI in health should protect the privacy of patients' information, ensuring their right to know which information about them is stored, to access their own medical records, and to have control over how their data should be used and disclosed, with a few exceptions.

The association emphasized that medical care provided by human beings should not be replaced by AI, but this technology has the potential to aid in providing these services, reducing system inefficiencies, and more adequately allocating resources. AI would be, ideally, a means to the objective of improving healthcare. Ama believes AI requires regulation, just like any other technology involved in diagnosing and treating patients. Government regulation of AI in healthcare should adequately protect patients, consumers, and health professionals, fostering trust in the system. Doctors would be responsible for ensuring that patients' health, well-being, and privacy are a priority over AI applications in healthcare[15].

The Australian association propositions seem reasonable, but, for a moment, one may question if such

a great responsibility must fall on doctors alone in case an AI systemic error applied to healthcare occurs. In an article investigating the rise of AI in medicine that used the qualitative analysis of documents and academic texts, Lamy and Malta[16] concluded that the ethical use of AI in healthcare also "depends on a practical and regulatory confrontation" of themes like prejudice and algorithmic discrimination (when data are biased), lack of transparency, intelligibility, and control.

In a similar direction, Dourado and Aith[17], after analyzing the General Personal Data Protection Law (LGPD) in the Brazilian legal system, write that

> Exercising the right to explanation in health depends on devising mechanisms to create explainable artificial intelligence systems and acknowledging the limits of algorithms' explicability. The scope of this right should be defined by criteria that must be elaborated by regulatory authorities and broadly discussed with society. This debate has only just begun.

Health professionals and researchers are overwhelmed by the speed with which new technological products are being introduced – another strategy that meets the market and technology's interest[7]. But does this interest the ethically responsible, economically consequential, and nationally secure science? These movements must be perceived from the simplest action, such as the ones that lead to the imbrication of free services offered by private companies and public administration of the three federative instances, that range from e-mails to "clouds" (expensive and energy-consuming computer servers), whose services may be suspended at any given time, as stated, by the way, on the terms of use of these blessings, even after a technological dependency has been established. Or precisely because of that? Unfortunately, it is not unusual to find defenders of the uncritical use of technology in the public sector using the argument filled with technological fetishism that one should "take the risk" of adopting new technologies, even when dealing with health. From these "details" emerge the range of technological products with no specific regulation for medicine or any other area, such as the case of generative AI – extraordinary, but not guaranteed to be fully secure, backed by the rhetoric power of marketing: *we have to* use them as they signify "development" and "progress". Thus, simple questions lead to complex answers: when something goes wrong, who is responsible? Who supervises it?

These doubts are not exclusive to public service. In a recent webinar, Girish N. Nadkarni, head of the Windreich Department of Artificial Intelligence and Human Health at the Icahn School of Medicine at Mount Sinai, a private entity located in New York, reached the same conclusion, in the healthcare sphere, found in the application of generative AI in the scope of social communication[7]: "We need to think a little bit around the risks surrounding it [AI use in medicine]. We deploy a lot of AI into clinical care, but I think that the decision not to deploy something is equally important as a decision to deploy it".[18]

Since these issues are not yet solved and there are no universally accepted methodologies, the Brazilian Internet Steering Committee simply recommends[19]: "The dialog between all interested actors who participate in the AI in healthcare ecosystem – especially developers, manufacturers, regulators, users, and patients – is highly advisable as the AI community grows". The Brazilian Internet Steering Committee also emphasizes the issue of transparency in building trust and communication strategies that meet the needs of end users (health professionals and managers), patients, and communities. Communication and transparency are thus interconnected with criteria that lead to adopting one technology instead of another, how it works, and what risks and benefits it can pose to the human body. This communication is so relevant *between* and *for* directly involved actors, like planners and society overall, that managers must learn *how to communicate* each situation.

If oncology needs to take responsibility for promoting this communication action, understanding it to be a geopolitical, multidisciplinary, and urgent action, then organized civil society must also be urged to understand that their safety cannot be delegated to the ethical thresholds of the AI market. Professors Andrew Barto and Richard Sutton, winners of the 2025 Turing Award, the "Oscar of Computing Science", for their technical contributions to "reinforcement learning", admitted being worried about the rush with which AI companies release new models without properly testing them – one of the consequences of the speed strategy referred earlier. To Rech *apud* Barto[20], "engineering practices evolved to try and mitigate negative consequences of technology, and I don't see it being practiced by the companies that develop it". That means, it is not actually the *technique* that can put people at risk in several ways with the use of AI in different domains, but *ethics*.

The ethics that directly or indirectly dialogue with people involved with the application of AI into healthcare will be tested at most between 2025-2030. In 2025, the share of the health AI market worldwide was estimated at around 31.25 billion dollars. In five years, this value would increase to 185.84 billion[21]. Many corporations will be interested in profiting the most with a double trend verified even in developing countries: 1. Population aging, a risk factor for many

diseases, including cancer, and mitigating the effects of climate change, which encourages more technology and; 2. Popularizing generative AI applications in several areas of knowledge, including healthcare, which will need to deal with the increase in infirmities.

This is a particularly delicate time as big techs have decided to relax their fact-checking systems, which, theoretically, also facilitates that stories on the harms and powers of technology spread through networks. Deregulation seems to interest businesses and normalizes technology acceptance without facing the regulatory rigors of other market segments – a legal exception to bits and qubits that is almost religious. China's progress, that recently surprised the world once again with the DeepSeek chatbot, whose training cost a fraction of its predecessors', only increased the urgency of USA tech companies, which should reinforce in everyone the need for extra care regarding AI when considering applying it to each area, especially in healthcare and in this intrinsically delicate political moment in the West.

So, while it may seem like good news that AI applications aimed at fighting cancer around the world are growing at an increasingly rapid pace (749 distinct records from 2015-2022), according to recent research by Fiocruz[22], the social universe with no specific legislation and active communication in which these technologies are deployed is also a cause for concern. Furthermore, the research confirms more signs of the technological race in oncology, with the USA surpassed by India and China in the number of patents. However, this race, healthy at first, should follow each country's legal and ethical compliance and be effectively communicated to society, not only in terms of results, but at each relevant step of the process. And, preferably, have more actors involved. Anyway, these gaps are verified in every nation to a greater or lesser extent.

The concrete fact is that few companies in a few countries and the hands of a few people dominate the AI market. If society does not acknowledge how big of a deal this is, nothing suggests this resource distribution will be different in healthcare, with predictable consequences for the less protected segments. This obviously includes oncology. In this specific case, there are many "business secrets" around AI investments, notably generative, in the oncology field. Many resources allocated to this technology are employed by private companies that rarely inform the public of the amount of invested capital. These are corporations such as Google, Microsoft, and IBM, in the USA, and research institutions, like the National Institutes of Health[23], that, in 2022, announced a multidisciplinary program of 130 million dollars in AI in healthcare, including cancer, over four years. The public institution made a point of highlighting its concern that the tools developed, and the data obtained do not perpetuate inequalities and observes ethical problems that may arise during the analysis and collection of information.

Unfortunately, these actions of transparency and proactivity are not observed in private companies and even public organizations worldwide. It is almost impossible to identify with certainty, item by item, the health segments that will receive resources for applications and research *with* and *in* AI. Which begs the question: why? What interests does this opacity, which does not seem to be a coincidence, but a project, serve?

CONCLUSION

Thus, it seems important to highlight that safety, transparency and social communication are fundamental regarding the implementation of generative or non-generative AI technologies in health applications and – due to all the stigmas it carries with it, social dramas and costs, in addition to the economic ones, in the order of billions[24] – even more so in cancer. Even if at this point in this argument limited by space there is still resistance to why clear and effective communication about the adoption of AI in healthcare is a necessarily multidisciplinary challenge in the short term, notably in oncology, consider that recent history has shown that there are no discursive voids: spaces not occupied by those who could have done something were quickly filled by different agents with often undisclosed interests. Such omissions have also impacted healthcare in budget cuts, denialist actions, and actions that discredit rational or scientific discourse, against human development and common security, especially in middle-income countries.

May History serve as an example, not an iterative error template.

CONTRIBUITION

The author participated in every stage of the manuscript drafting and approved the final version to be published.

DECLARATION OF CONFLICT OF INTERESTS

There is no conflict of interest to declare.

FUNDING SOURCES

None.

REFERENCES

1. Check Point [Internet]. [Redwood]: Check Point Software Technologies; ©1994-2024. The state of cyber security 2024, 2024 [Acesso 2025 fev 21]. Disponível em: https://engage.checkpoint.com/security-report-2025/items/report--cyber-security-report-2024?utm_campaign=dg-cm_ps_25q1_ww_all_mix-gen-inf-cyber-security-report-2025-or_en-gsem&utm_source=google-dg&utm_medium=cpc&utm_term=cybersecurity±reports±2024&utm_content=cs_report&gad_source=1&gclid=CjwKCAiAn9a9BhBtEiwAbKg6fu7YOa0qbO5PzhR8i6r3cXKwm0GrawIbMpWCsABaO3betuc0h3RLFRoCqWUQAvD_BwE
2. Check Point [Internet]. [Redwood]: Check Point Software Technologies; 2025. The state of cyber security 2025, 2025 [Acesso 2025 fev 19]. Disponível em: https://engage.checkpoint.com/security-report-2025?utm_campaign=dg-cm_ps_25q1_ww_all_mix-gen-inf-cyber-security-report-2025-or_en-gsem&utm_source=google-dg&utm_medium=cpc&utm_term=state%20of%20cybersecurity%202025&utm_content=state_cybersec&gad_source=1&gclid=CjwKCAiAn9a9BhBtEiwAbKg6frHixQ--XsFgDubZop5p7WRiNiIP99dPO_Kcjd4FlQUaK5ZoqeLzlRoCGMgQAvD_BwE
3. Reuters [Internet]. Londres: Reuters; [sem data]. Google desenvolve co-cientista de IA para auxiliar pesquisadores, 2025 fev 19. [Acesso 2025 fev 21, Atualizado em 2025 abr 19]. Disponível em: https://www.reuters.com/technology/artificial-intelligence/google-develops-ai-co-scientist-aid-researchers-2025-02-19/
4. Conselho Federal de Biomedicina [Internet]. Brasília, DF: CFBM; [sem data]. O que faz o Biomédico, 2024. [Acesso 2025 fev 21]. Disponível em: https://cfbm.gov.br/profissionais/o-que-faz-o-biomedico/#:~:text=O%20Biom%C3%A9dico%20%C3%A9%20responsável%20pela,produz%20vacinas%20para%20combat%C3%AA%2Dlas
5. Angwin J, Larson J, Mattu S. Machine Bias: there's software used across the country to predict future criminals and it's biased against blacks. Pro Publica [Internet]. 2016 maio 23. [Acesso 2025 fev 9]. Disponível em: https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
6. Gerken T. AI cracks superbug problem in two days that took scientists Years. BBC [Internet]. 2025 fev 20. [Acesso 2025 fev. 24]. Disponível em: https://www.bbc.com/news/articles/clyz6e9edy3o
7. Amaral Filho N. Jornalismo e IA generativa: ensino, interdisciplinaridade, mercado e ética. Rio de Janeiro: Mauad X; 2024. 180p.
8. Lee KF. Inteligência artificial: como os robôs estão mudando o mundo, a forma como amamos, nos relacionamos, trabalhamos e vivemos. Rio de Janeiro: Globo Livros; 2019.
9. Google Cloud [Internet]. [local desconhecido]: Google Cloud; [sem data]. Tendências de IA para negócios em 2025. [Acesso 2025 fev 28]. Disponível em: https://cloud.google.com/resources/ai-trends-report?hl=pt-BR
10. Liu X, Faes L, Kale AU, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digit Health. 2019;1(6):271-97. doi: https://doi.org/10.1016/s2589-7500(19)30123-2
11. Ayers JW, Poliak A, Dredze M. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. JAMA Intern Med. 2023;183(6):589-96. doi: https://doi.org/10.1001/jamainternmed.2023.1838
12. Instituto Nacional de Câncer. Rede Câncer [Internet]. 2024 [Acesso 2025 fev 28];56:1-44. Disponível em: https://ninho.inca.gov.br/jspui/handle/123456789/17180
13. Garg RK, Urs VL, Agarwal AA, et al. Exploring the role of ChatGPT in patient care (diagnosis and treatment) and medical research: a systematic review. Health Promot Perspect. 2023;13(3):183-91. doi: https://doi.org/10.34172/hpp.2023.22
14. Gormley B. This startup wants to use AI to uncover dangerous drugs. The Wall Street Journal [Internet]. 2025 fev 12. [Acesso 2025 mar 7]. Disponível em: https://www.wsj.com/articles/this-startup-wants-to-use-ai-to-uncover-dangerous-drugs-40297123
15. Australian Medical Association [Internet]. [local desconhecido]: AMA; 2023. Position Statement. Artificial Intelligence in Healthcare, 2023. [Acesso 2025 fev 25]. Disponível em: https://www.ama.com.au/sites/default/files/2023-08/Artificial%20Intelligence%20in%20Healthcare%20-%20AMA.pdf
16. Lamy M, Malta KC. Avanços e riscos da inteligência artificial na atenção à saúde. LSS [Internet]. 2023 [Acesso 2025 fev 26];12(2):108-19. Disponível em: https://periodicos.unisanta.br/LSS/article/view/936
17. Dourado DA, Aith FMA. A regulação da inteligência artificial na saúde no Brasil começa com a Lei Geral de Proteção de Dados Pessoais. RSP. 2022;56:1-7. doi: https://doi.org/10.11606/s1518-8787.2022056004461
18. Fiddler [Internet]. [local desconhecido]: Fiddler; [sem data]. AI explained: genAI use cases and challenges in healthcare. [Acesso 2025 mar 03]. Disponível em: https://www.fiddler.ai/webinars/ai-explained-genai-use-cases-and-challenges-in-healthcare?utm_source=mkto&utm_medium=email&utm_campaign=genai-healthcare&mkt_tok=NTEzLVJQUS02OTkAAAGY7LDtzKLF0IbPJTIS_ffwXHVsUElPQjrfsKd5oo6LRrWEDuevAmZrhDW_Y1fWCMI7ECmjhIaF81cIDx9JlroM1-Pp-ShT0osaW8czSCnpYw

19. Comitê Gestor da Internet no Brasil. Inteligência artificial na saúde: potencialidades, riscos e perspectivas para o Brasil [Internet]. São Paulo: NIC.br; 2024. [acesso 2025 fev. 15]. 338p. Disponível em: https://cetic.br/media/docs/publicacoes/7/20240903150639/estudos_setoriais-ia-na-saude.pdf
20. Rech R. Vencedores do prêmio turing fazem alerta sobre IAs sem sistemas de segurança. Exame [Internet]. 2025 mar 5. [Acesso 2025 mar. 7]. Disponível em: https://exame.com/inteligencia-artificial/vencedores-do-premio-turing-fazem-alerta-sobre-ias-sem-sistemas-de-seguranca/
21. Mordor Intelligence [Internet]. [local desconhecido]: Mordor Intelligence; [sem data]. Artificial intelligence in medicine market size & share analysis - growth trends & forecasts (2025 - 2030). [Acesso 2025 mar 10]. Disponível em: https://www.mordorintelligence.com/industry-reports/artificial-intelligence-in-medicine-market
22. Braga L, Lopes R, Alves L, et al. The global patent landscape of artificial intelligence applications for cancer. Nat Biotechnol. 2023;41(12):1679-87. doi: https://doi.org/10.1038/s41587-023-02051-9
23. National Institute of Health [Internet]. [local desconhecido]: NIH; 2022. NIH launches Bridge2AI program to expand the use of artificial intelligence in biomedical and behavioral research, 2022 set 13. [Acesso 2025 mar 11]. Disponível em: https://www.nih.gov/news-events/news-releases/nih-launches-bridge2ai-program-expand-use-artificial-intelligence-biomedical-behavioral-research
24. Instituto Nacional de Câncer [Internet]. Rio de Janeiro: INCA; 2022. Sumário executivo: gastos federais atuais e futuros com os cânceres atribuíveis aos fatores de risco relacionados à alimentação, nutrição e atividade física no Brasil, 2022. [Acesso 2025 mar 13]. Disponível em: https://www.inca.gov.br/publicacoes/relatorios/gastos-federais-atuais-e-futuros-com-os-canceres-atribuiveis-aos-fatores-de

Recebido em 17/3/2025
Aprovado em 17/3/2025

Scientific-editor: Anke Bergmann. Orcid iD: https://orcid.org/0000-0002-1972-8777