RESEARCH ETHICS
J. Heal Ethics Admin
Volume 11 | Number 4 (Fall 2025)
www.jheaonline.org
ISSN 2474-2309 | doi:10.22461/jhea.1.71653
https://doi.org/10.22461/jhea.1.71653
PublishedNov 04, 2025

BIOETHICS AND ARTIFICIAL INTELLIGENCE IN MEDICINE: A CRITICAL-ARGUMENTATIVE APPROACH

Evangelina Mollar, MBE[1], Mariano Sorrentino, Dr.[2]
emollar@bioethx.org

Abstract: *This article explores the growing role of Artificial Intelligence (AI) in medicine, with a special focus on the surgical field, and analyzes the profound bioethical dilemmas that its implementation raises. While AI offers significant advances, such as high-precision robotic surgery, standardized evaluation of professional performance, and optimization of clinical processes, its adoption requires critical reflection to avoid the dehumanization of medical practice.*

The text addresses several key aspects:

Medical Training: It examines the paradigm shift in surgical education, where AI, through the analysis of surgical videos, can offer immediate feedback that is less variable than some human ratings, in contrast to the traditional model of human mentorship. Although this technology is efficient, it raises the question of what is lost by replacing the inspiring figure of the human mentor who transmits values, empathy, and emotional support.
Surveillance and Autonomy: The use of AI-based video surveillance systems in operating rooms raises questions about professional privacy and autonomy. The article warns about the "Hawthorne effect," where constant monitoring could increase stress and pressure on professionals, rather than fostering an environment of continuous improvement.

Justice and Equity: It discusses the inequalities that AI can exacerbate. The asymmetrical access to these technologies between countries and between the public and private sectors, coupled with algorithmic biases trained on data from specific populations, could deepen healthcare gaps.

*Address correspondence to: Evangelina Mollar. Lawyer, Bioethics Specialist. Email: emollar@bioethx.org

+To cite this article: Mollar, E., Sorrentino, M. “Bioethics and Artificial Intelligence in Medicine: A Critical-Argumentative Approach ”. The Journal of Healthcare Ethics & Administration Vol. 11, no. 4 (Fall 2025): 38-47, https://doi.org/10.22461/jhea.1.71653

[1] Lawyer, specialist in Bioethics, Member of the Bioethics Committee of INCUCAI, Regional Director for Latin America of Bioethx
[2] General surgeon, specialist in head and neck surgery, microsurgery, and transoral robotic surgery

Governance and Responsibility: Finally, the article concludes that the implementation of medical AI is not a neutral process. It advocates for shared and transparent ethical governance, where clear responsibilities are defined between humans and machines. It proposes an "anticipatory bioethics" that, instead of reacting to problems, prevents them by integrating ethical principles from the design phase of the technology. The ultimate goal is to use AI as an ally of care, not as a substitute for the human dimension of medicine, ensuring that technological advancement is oriented toward social well-being and human dignity.

Keywords: *Artificial Intelligence (AI), Medical AI, Robotic Surgery, Bioethics, Surgical Education, Professional Autonomy, Algorithmic Bias, Healthcare Equity, Ethical Governance, Human Dignity.*

I. INTRODUCTION

In recent decades, Artificial Intelligence (AI) has moved from being a futuristic promise to a tangible tool that is increasingly present in many areas of human life. Medicine, as a science and art profoundly influenced by technological development, has not been immune to this transformation. The progressive incorporation of AI systems in diagnosis, treatment, and clinical management represents one of the most significant changes of the 21st century in the field of health. In surgical practice, in particular, the use of AI has become increasingly relevant, not only as a complementary tool in minimally invasive surgeries but also as a resource for the standardized evaluation of professional performance (Topol, 2019).

However, the enthusiasm for technological innovation coexists with a set of bioethical questions that demand rigorous analysis. What impact does AI have on medical training? How does it affect the doctor-patient relationship and professional autonomy? Does efficiency justify the partial replacement of human figures like the surgical mentor? What risks exist regarding privacy, justice, and dehumanization? These questions are fundamental to evaluating the role that technology should play in medicine, a field that continues to champion its humanistic vocation (Moreno, 2021).

In the last decade, we have witnessed an exponential growth in the incorporation of advanced technologies in medicine, particularly in the surgical field. Robotic surgery, in particular, has become increasingly relevant in the context of minimally invasive surgery, complementing—and in some cases replacing—conventional laparoscopic techniques. This advance has led to a paradigm shift not only in how we operate but also in how we train new generations of surgeons.

In the American academic model, the surgical training system is based on the traditional mentorship format, where an experienced surgeon guides and provides real-time feedback to residents, modeling not only technical skills but also professional attitudes. However, at our hospital, we have recently initiated a clinical trial that evaluates a disruptive alternative: artificial intelligence (AI) as a surgical evaluation tool.

This study exemplifies a specific use of AI: residents are randomly assigned to two groups. The first group continues under the traditional human mentorship system. The second group is evaluated by AI, which analyzes videos of minimally invasive surgeries performed by the residents. These recordings are compared with a vast database of interventions performed by expert surgeons. The algorithm breaks down each procedure into multiple metrics: unnecessary movements, inefficient use of instruments, operating time, movement efficiency, among others. Based on this, it generates standardized and auditable reports with specific suggestions on which aspects of performance to improve.

This approach, while technologically impressive, raises significant bioethical and pedagogical questions. AI lacks the human dimension of the surgical mentor—that inspiring figure who forges bonds, transmits values, and provides vocational motivation. However, it offers notable advantages: immediate, continuous, and standardized feedback. What do we lose and what do we gain when we allow an algorithm to partially complement the human mentor?

Another current example is the use of permanent video surveillance systems in operating rooms. Multiple cameras record the actions of the surgical team in real-time. AI analyzes this data to evaluate efficiency in aspects such as patient preparation, the transition between cases, and unjustified delays. Here another dilemma arises: the possibility of the Hawthorne effect, whereby individuals modify their behavior simply by knowing they are being observed. Are we promoting an environment of continuous improvement or one of permanent surveillance?

AI, like any powerful tool, has immense potential to elevate the quality of medical care and optimize professional performance. However, it can also displace essential human dimensions, such as empathy, inspiration, and the vocation for service. In an already emotionally demanding profession, the incorporation of new technologies should not become an additional burden but rather a resource that respects the mental health of professionals and the fundamental values of medicine.

The real challenge will be to find an ethical balance between efficiency and humanity. We must ensure that AI complements, without replacing, the essence of medicine as practiced by people who chose this path driven by a genuine spirit of service.

II. THE RISE OF ARTIFICIAL INTELLIGENCE IN SURGERY AND CLINICAL MEDICINE

Artificial intelligence (AI) has become one of the most promising tools for the future of medicine. Its ability to analyze large volumes of data in real time, identify complex patterns, and optimize clinical processes is transforming not only how diseases are diagnosed and treated, but also how healthcare systems are organized and how medical professionals are trained. In the surgical field, these advances have taken on special relevance, promoting a silent but profound transformation of the medical act and its many dimensions.

Among the most notable applications is robot-assisted surgery, which allows for high-precision, minimally invasive procedures. The Da Vinci system, for example, is one of the most widespread globally and has been adopted in high-level hospitals in developed countries as well as in some centers in Latin America. This type of technology improves surgeon ergonomics, reduces physiological tremors, and reaches complex anatomical areas with greater precision than traditional laparoscopic surgery (Herron et al., 2020).

Along with these tools, AI has begun to be used in the automatic analysis of medical images—X-rays, MRIs, CT scans—in some cases surpassing human diagnostic accuracy. It is also used in clinical decision support systems, recommending personalized treatments based on the analysis of databases of millions of patients, as well as in the management of hospital resources, optimizing schedules, beds, and surgical times (Jiang et al., 2017).

In the field of medical training, this study exemplifies a specific and novel use of AI: the evaluation of surgical performance based on video analysis. This approach represents a paradigm shift in medical education by introducing standardized and auditable criteria to assess clinical skills, reduce the variability of traditional evaluation, and offer immediate feedback to residents. However, as will be discussed in the following chapters, this model is not free of ethical tensions, particularly regarding the humanization of learning and the emotional support of the future professional.

Another emerging field is the use of AI systems in anesthesiology and intraoperative monitoring. Algorithms capable of anticipating adverse events, detecting arrhythmias in real time, or automatically adjusting drug administration are beginning to change the logic of the operating room, with potential benefits for patient safety but also with new bioethical challenges related to responsibility, professional autonomy, and technology oversight.

The technological acceleration has been so significant that even major international health organizations have begun to publish specific guidelines on ethics and AI. In 2021, the World Health Organization (WHO) warned that while these technologies can improve health equity, they can also reproduce inequalities, violate rights, and alter the relationship between professionals and patients if not properly regulated and evaluated (World Health Organization, 2021).

In this context, it is essential that the enthusiasm for technological innovation be accompanied by critical and multidisciplinary reflection. Medicine is not just about technique or efficiency; it is also about human relationships, prudent judgment, and comprehensive care for the person. The challenge is not to stop technological advancement but to integrate it in an ethical, just, and humanized way.

III. THE TRANSFORMATION OF MEDICAL TRAINING: HUMAN MENTOR OR ALGORITHM?

Historically, medical training has been conceived as an intensely human process. Knowledge is transmitted not only through technical content but also through direct experience, observing by example, and building pedagogical bonds. The figure of the mentor—an experienced doctor who guides, evaluates, orients, and motivates students and residents—has been central to this process. In the surgical field, this relationship is even more valuable, as the acquisition of technical skills requires close supervision, continuous practice, and individualized feedback (Cruess et al., 2008).

However, the irruption of artificial intelligence (AI) into training processes is beginning to challenge traditional teaching models. As previously described, the use of algorithms capable of analyzing videos of surgical interventions allows for standardized and auditable evaluation of residents' performance, breaking down each procedure into variables such as operative time, inefficient movements, or inappropriate use of instruments. Based on this data, personalized reports are generated that offer specific suggestions for improving technique.

This automated evaluation system has undeniable advantages: it reduces personal bias, allows for immediate feedback, and standardizes criteria that, in the past, could vary depending on the tutor's perception. Additionally, it can be used as a complementary resource to human training, enriching education with empirical data and comparative metrics. Nevertheless, its implementation also raises significant bioethical and pedagogical questions.

First, algorithmic evaluation does not account for subjective but fundamental dimensions of medical practice, such as decision-making in contexts of uncertainty, empathy, clinical judgment, and communication with the patient and team. These soft skills—often invisible to quantitative data—are essential for the formation of an integral professional, capable of practicing with sensitivity and responsibility (Hodges & Lingard, 2012).

Second, there is a risk of blurring the figure of the surgical mentor, who not only fulfills a teaching role but also a formative one in an ethical and vocational sense. The mentor represents a professional role model who transmits values, commitment, and emotional support. The total or partial substitution of this bond with an algorithmic evaluation could impoverish the educational experience, weakening the construction of a professional identity in future doctors (Wear & Aultman, 2005).

From a bioethical perspective, this dilemma creates tension between two fundamental principles: beneficence and autonomy. On one hand, it is necessary to guarantee quality training based on standardized data that allows for improvement of surgical technique. On the other, the comprehensive development of the person in training must be respected, preventing technology from reducing their experience to a set of standardized, auditable, or less variable than human rating metrics. Furthermore, the principle of justice also comes into play: access to these technologies is not universally guaranteed, which could create gaps between institutions with more or fewer resources.

IV. SURVEILLANCE, TRANSPARENCY, AND PRIVACY IN SURGICAL ENVIRONMENTS

Artificial intelligence (AI) has not only revolutionized surgical practice in terms of technical precision and professional performance evaluation but has also introduced new dynamics of institutional surveillance and control, with significant bioethical implications. The progressive installation of permanent video surveillance systems in operating rooms, connected to algorithms that analyze the activity of the medical team in real time, raises

fundamental questions about privacy, autonomy, professional freedom, and the workplace climate within health centers.

These systems typically record multiple variables: preparation times, surgery duration, transitions between procedures, interruptions, technical errors, and overall team efficiency. In principle, this data can be valuable for identifying bottlenecks, improving patient safety, and promoting quality processes (Sheikh et al., 2021). However, when applied in clinical contexts where human decisions and actions have unpredictable margins, the presence of constant surveillance can generate unintended side effects.

One of the most significant is the so-called Hawthorne effect, where people modify their behavior when they know they are being observed. In high-stakes surgical environments, this additional pressure can interfere with concentration, increase stress, and harm the emotional well-being of the team, particularly residents and professionals in training (McCambridge et al., 2014). Instead of fostering an environment of continuous improvement, algorithmic surveillance can establish a culture of evaluative pressure, where intrinsic motivation is influenced by external metrics.

From a bioethical perspective, these dynamics directly affect the principle of professional autonomy, as surgeons may see their ability to make decisions limited by the fear of algorithmic interpretations. Additionally, the principle of non-maleficence is relevant, as a hyper-surveilled environment can negatively affect the mental health of healthcare workers in an already emotionally demanding profession.

Furthermore, professional privacy is also challenged. Although audiovisual records are often presented as being for academic or quality control purposes only, in practice, there are sometimes unclear regulatory frameworks on who accesses this information, for what purposes, for how long, and under what data protection measures for both personal and professional data. This ambiguity can generate mistrust among both professionals and patients, whose procedures are also being recorded.

In addition, the unequal implementation of these technologies between highly complex private centers and public hospitals introduces a dilemma of distributive justice. While some professionals have AI as support to improve their practice, others may be subjected to control mechanisms without symmetrical access to the educational benefits or the right to challenge automated reports.

From a critical perspective, it is essential that the incorporation of surgical video surveillance and analytical AI systems be subject to ethical governance processes, which should include:

Informed consent (professional and, if applicable, patient),
Access and confidentiality policies,
Participation of bioethics committees in the review of protocols,
Human supervision of reports generated by AI, and
Mechanisms for emotional and legal protection against algorithmic errors or failures.

As the WHO points out in its report on ethics and artificial intelligence, the protection of human dignity and fundamental rights must be a priority in all processes of technological innovation in health (World Health Organization, 2021). The goal is not to oppose surveillance per se, but to prevent the promise of efficiency from becoming a threat to the respect, freedom, and well-being of medical staff.

V. JUSTICE AND EQUITY IN ACCESS TO MEDICAL ARTIFICIAL INTELLIGENCE

The implementation of AI-based technologies in healthcare has been seen as a promising advance for early diagnosis, personalized treatment, reduced errors, and optimized clinical resources. However, this technical progress has not been, nor will it be, universal or equitable. In healthcare contexts marked by structural inequality, the incorporation of AI does not escape the bioethical tensions generated by the asymmetrical distribution of technological benefits, particularly in countries in the Global South.

From a bioethical perspective focused on distributive justice, access to these innovations raises crucial questions: Who actually gets access to medical AI? Who is excluded? What kind of care do people who are not part of AI-enabled systems receive? Are we amplifying existing gaps between public and private, rural and urban, central and peripheral health systems?

In Latin America, for example, health systems have multiple limitations: deficient infrastructure, a lack of digital interoperability, poor connectivity, a shortage of trained personnel, and fragmented budgets. In this context, the introduction of AI in high-complexity centers or elite private clinics can become a form of technological elitism, leaving behind vulnerable populations who continue to receive care under traditional, less efficient, and even obsolete care models (Bertuzzi & Luna, 2023).

Furthermore, this inequality is not just geographical or socioeconomic; it's also epidemiological. AI systems trained with data from European or North American populations often have algorithmic biases when applied to diverse contexts. This can result in diagnostic errors or inadequate treatments for ethnically distinct groups with unique clinical profiles (Obermeyer et al., 2019). This situation violates the principle of epistemic justice, which requires considering human diversity in the development of health tools.

On the other hand, there is a problem of technological and data sovereignty. Many countries lack the technical and regulatory capacities to design their own AI systems, depending on imported technologies developed by transnational corporations. This raises bioethical risks related to the exploitation of sensitive data, a lack of control over automated decision-making processes, and the subordination of local health systems to commercial interests that are external to the common good (Floridi et al., 2018).

A bioethics committed to equity therefore demands the formulation of technological integration policies that promote social justice and digital inclusion. This involves:

Investing in infrastructure and connectivity for vulnerable areas,
Creating regulatory frameworks that protect patient rights and ensure transparency,
Fostering the interoperability of public systems,
Developing locally trained algorithms with diverse data, and
Strengthening local capacities for innovation and ethical control.

As UNESCO (2021) emphasizes, AI in health should be seen as a global common good, whose distribution is not dictated by market logics, but guided by the principles of equity, solidarity, and human dignity. The ethics of justice demand not only expanding access to these technologies but also ensuring that their implementation does not reproduce or deepen the injustices already present in our healthcare systems.

VI. ETHICAL GOVERNANCE AND SHARED RESPONSIBILITY IN THE DEVELOPMENT OF MEDICAL ARTIFICIAL INTELLIGENCE

The development and implementation of AI in medicine is not an exclusively technical or neutral process. It involves decisions with profound ethical, legal, and social implications. Because of this, it is necessary to build robust ethical governance that can guide the design, use, and evaluation of these technologies based on the common good, respect for human dignity, and the protection of patients' fundamental rights.

The ethical governance of medical AI should be understood as a framework of actors, principles, regulations, and institutional mechanisms that collectively regulate how these tools are designed and deployed. Unlike a purely technical or self-regulated model, ethical governance requires the active participation of multiple sectors: governments, healthcare professionals, developers, patients, bioethics committees, academics, and civil society (Morley et al., 2020).

One of the main challenges is to establish clear responsibilities in contexts where decision-making is mediated by algorithms. Who is responsible if an AI system makes a diagnostic error that results in harm to a

patient? The programmer? The medical institution that uses it? The doctor who follows its recommendations? These questions require a redefinition of professional responsibility to include forms of co-responsibility between humans and machines (Topol, 2019).

In addition, it is essential to promote algorithmic transparency. Many current AI systems function as "black boxes," whose internal processes are not understandable or auditable by users. This opacity risks the principle of patient autonomy, as patients should be able to make informed decisions about their health. It also hinders public control and accountability, two essential conditions for any ethically acceptable technology (Floridi et al., 2018).

In the face of these risks, various international organizations have proposed governance frameworks for an ethical AI. For example, the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021) establishes guiding principles such as inclusion, sustainability, gender equity, respect for human rights, and the promotion of human and environmental development. The European Commission has also published guidelines for a trustworthy AI, focusing on seven key requirements: human supervision, technical robustness, privacy, transparency, diversity, social well-being, and accountability (European Commission, 2019).

In this context, hospital bioethics committees take on a fundamental role as mediating and evaluating bodies. They should not only act in the face of specific dilemmas but also serve as spaces for proactive analysis, advice, and monitoring of the impact these technologies may have on the doctor-patient relationship, access equity, and quality of care. The promotion of anticipatory ethics—which does not react to conflicts after they occur but foresees them—is key in highly technologized environments.

Finally, a cross-cutting bioethical education is required in medical and engineering training programs. AI developers must know the principles of bioethics, just as healthcare professionals must understand the limitations and risks of the algorithmic systems they use. Ethical governance is, ultimately, a shared task: it does not depend on a single actor but on an ethical community willing to think critically about the technological direction of contemporary medicine.

VII. CONCLUSIONS AND PROPOSALS FOR A BIOETHICS OF ARTIFICIAL INTELLIGENCE IN MEDICINE

The incorporation of AI in medicine represents one of the greatest ethical challenges of the 21st century. Far from being a neutral tool, AI in health is transforming how we diagnose, treat, research, and make clinical decisions, with direct consequences on patient autonomy, the equity of the healthcare system, the doctor-patient relationship, and professional responsibility.

Throughout this work, substantial risks have been identified: from algorithmic biases and system opacity to tensions in confidentiality, distributive justice, and human supervision. The potential of these technologies to improve access, efficiency, and medical precision has also been highlighted, as long as they are framed within an ethics of care, oriented toward respect for human dignity.

In this sense, a bioethics of AI in medicine cannot be limited to a reactive analysis of harm. Instead, it must promote anticipatory ethics, capable of intervening in the initial stages of technological design and development. For this, an interdisciplinary approach is essential, one that integrates healthcare professionals, engineers, ethics experts, legislators, patients, and representatives of civil society.

The main proposals arising from this analysis are:

Incorporate bioethical training into medical, engineering, and informatics programs, so that those who design algorithms understand the moral dilemmas associated with their implementation in clinical contexts.

Foster institutional instances of permanent bioethical evaluation, such as interdisciplinary committees that advise on the adoption of algorithmic technologies, supervise their operation, and analyze their social impacts.

Promote national and international regulatory frameworks that ensure transparency, accountability, and human supervision, as proposed by UNESCO (2021) and the European Commission (2019), among others.

Guarantee the participation of patients and civil society in decision-making about the use of AI in health, reinforcing principles of justice, equity, and digital inclusion (Floridi et al., 2018).

Revise the classic notion of medical responsibility, adapting it to a context in which clinical decisions may be influenced by algorithmic systems that are not completely comprehensible to professionals.

Establish shared ethical governance policies, in which technological development is guided by the common good, public interest, and the promotion of person-centered medicine (Sorrentino, 2023; Topol, 2019).

Personalist bioethics, with its focus on the inviolable dignity of the human being and the relational nature of care, offers a solid theoretical framework to guide the responsible development and use of medical AI. This approach allows us to transcend purely utilitarian models and technocratic visions, proposing a balance between technological innovation, healthcare justice, and respect for human life.

Ultimately, the future of AI in medicine will depend not only on scientific advances but-fundamentally-on the ethical decisions we make as a global community. An informed, plural, and anticipatory bioethical commitment will be key to ensuring that these technologies become allies of care and not new sources of exclusion or dehumanization.

REFERENCES

Beauchamp, T. L., & Childress, J. F. (2019). Principles of biomedical ethics (8th ed.). Oxford University Press.

Bertuzzi, L., & Luna, F. (2023). Inequality and justice in digital medicine: Ethical challenges for Latin America. Revista Latinoamericana de Bioética, 23(1), 45– 62.

Cruess, S. R., Johnston, S., & Cruess, R. L. (2008). Professionalism for medicine: Opportunities and obligations. Medical Journal of Australia, 189(8), 368–371. https://doi.org/10.5694/j.1326-5377.2008.tb02102.x

European Commission. (2019). Ethics guidelines for trustworthy AI. Independent High-Level Expert Group on Artificial Intelligence.

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and Machines, 28(4), 689–707. https://doi.org/10.1007/s11023-018-9482-5

Herron, D. M., Marohn, M., & Efron, J. (2020). Robotic surgery in general surgery residency: A review of the current state and future directions. Journal of Surgical Education, 77(4), 837–843. https://doi.org/10.1016/j.jsurg.2019.12.009

Hodges, B. D., & Lingard, L. (Eds.). (2012). The question of competence: Reconsidering medical education in the twenty-first century. Cornell University Press.

Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... & Wang, Y. (2017). Artificial intelligence in healthcare: Past, present and future. Stroke and Vascular Neurology, 2(4), 230–243. https://doi.org/10.1136/svn-2017-000101

McCambridge, J., Witton, J., & Elbourne, D. R. (2014). Systematic review of the Hawthorne effect: New concepts are needed to study research participation effects. Journal of Clinical Epidemiology, 67(3), 267–277. https://doi.org/10.1016/j.jclinepi.2013.08.015

Moreno, J. D. (2021). Ethics and the future of artificial intelligence in medicine. Hastings Center Report, 51(1), 10–12. https://doi.org/10.1002/hast.1206

Morley, J., Machado, C. C. V., Burr, C., Cowls, J., Joshi, I., Taddeo, M., & Floridi, L. (2020). The ethics of AI in health care: A mapping review. Social Science & Medicine, 260, 113172. https://doi.org/10.1016/j.socscimed.2020.113172

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453. https://doi.org/10.1126/science.aax2342

Sheikh, A., Anderson, M., Albeyatti, A., & Mossialos, E. (2021). Harnessing the potential of artificial intelligence for learning health systems: A critical review and proposed framework. Health Affairs, 40(9), 1376–1384.

Topol, E. (2019). Deep medicine: How artificial intelligence can make healthcare human again. Basic Books.

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. https://unesdoc.unesco.org/ark:/48223/pf0000381137

Wear, D., & Aultman, J. M. (2005). The limits of narrative: Medical student resistance to confronting inequality and oppression in literature and beyond. Medical Education, 39(10), 1056–1065. https://doi.org/10.1111/j.1365-2929.2005.02274.x

World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200