HEALTHCARE UNDER THE EU AI ACT: RISKS, RIGHTS, AND RESPONSIBILITY IN EUROPEAN FRAMEWORK

SAMET TATAR

HEALTHCARE UNDER THE EU AI ACT: RISKS, RIGHTS, AND RESPONSIBILITY IN EUROPEAN FRAMEWORK

Samet TATAR
Ankara Yıldırım Beyazıt University

Abstract

Artificial intelligence (AI) is transforming healthcare in Europe, offering unprecedented opportunities for diagnosis, treatment, and patient monitoring while raising profound legal and ethical challenges. The European Union (EU)'s Artificial Intelligence Act (AI Act), enacted in 2024, represents the world's first comprehensive regulatory framework for AI and establishes a risk-based approach that automatically classifies all healthcare applications of AI as "high-risk." This article provides an in-depth analysis of the AI Act's implications for the healthcare sector, situating it within the EU's broader tradition of rights-based digital governance and examining its interaction with existing regimes such as the Medical Device Regulation (MDR) and the In Vitro Diagnostic Regulation (IVDR). The discussion explores how the AI Act redefines the responsibilities of developers, providers, healthcare institutions, and professionals through obligations of transparency, data governance, documentation, and human oversight. It also evaluates the ethical and human rights dimensions of the Act, linking provisions on bias, opacity, and accountability to the Charter of Fundamental Rights of the European Union (CFEU). Special attention is paid to the compliance burdens facing small and medium-sized enterprises (SMEs), which often drive innovation in health technology but face disproportionate challenges under the current regulatory model.

Keywords

EU AI Act, healthcare AI, risk-based regulation, GDPR, MDR/IVDR, fundamental rights, human oversight

1. INTRODUCTION

AI has rapidly moved from research laboratories and dystopian speculation into hospital corridors across Europe, transforming the ways in which medicine is practiced and healthcare is delivered. Radiology departments increasingly rely on AI to detect anomalies in thousands of images, predictive analytics assist clinicians in identifying high-risk patients before symptoms emerge, and intensive care units benefit from algorithmic alerts that anticipate sudden deterioration. The COVID-19 pandemic accelerated these developments, as algorithms were used to model infection waves, allocate scarce resources, and speed vaccine discovery. These examples reflect both the transformative potential of AI and the unease it generates when applied to matters as personal and consequential as healthcare (Busch et al., 2024, p. 1-2; RAND Europe, 2024; Mello & Cohen, 2025).

Real-world experiences illustrate this tension. IBM's Watson for Oncology was once promoted as a revolutionary tool capable of matching patient data with global clinical trials, yet it soon faced criticism for overpromising and producing unreliable recommendations. Google DeepMind's collaboration with the UK National Health Service produced remarkable advances in diagnostic imaging but also drew backlash when millions of patient records were shared without sufficient consent safeguards. Such cases underscore the dual reality of healthcare AI: it can inspire hope for faster, more accurate care, while simultaneously eroding public trust if ethical and legal safeguards are absent (Goisauf et al., p. 2-3, 2025; Mello & Cohen, 2025).

The legal and ethical challenges are therefore significant. Many AI systems operate as "black boxes," producing recommendations through processes too complex for meaningful explanation. When outcomes are favorable, opacity may be tolerated; when harm occurs, it creates uncertainty about responsibility. Was the fault with the developer who trained the system, the institution that deployed it, the doctor who relied on it, or the regulator who failed to prevent it? This accountability gap reveals the inadequacy of traditional liability frameworks when applied to algorithmic medicine. Bias in datasets compounds these problems. Medical algorithms trained on narrow datasets perform less effectively for women, minorities, or patients outside Europe. At the same time, the massive datasets fueling AI raise pressing privacy concerns, particularly where lines blur between medical research and commercial exploitation. Human dignity itself is at stake when life-altering medical decisions are delegated to opaque computational logic beyond patients' comprehension (Price, 2024, p. 150; Solaiman, 2025 p. 3-5).

In response to these issues, the EU enacted the AI Act. This groundbreaking legislation establishes the world's first comprehensive framework for governing AI. Building on earlier digital regulation, most notably the General Data Protection Regulation (GDPR), the AI Act employs a risk-based model: systems deemed to create "unacceptable risk" are banned outright

under Article 5, while those classified as "high-risk" must comply with rigorous requirements concerning data governance (Art. 10), technical documentation (Art. 11), logging (Art. 12), transparency (Art. 13), human oversight (Art. 14), and accuracy, robustness, and cybersecurity (Art. 15). All AI systems deployed in healthcare- whether for diagnosis, treatment, or clinical risk prediction- are automatically classified as high-risk under Annex III. This classification underscores the life-and-death stakes of medical decisions and the EU's prioritisation of patient safety, while also signalling the Union's broader commitment to fairness, privacy, and dignity.

Yet automatic classification carries costs. Compliance with the AI Act's strict requirements may be feasible for multinational corporations, but it risks overwhelming SMEs, which form the backbone of health technology innovation in Europe. This raises a fundamental policy dilemma: how can regulators protect patients without stifling innovation? The challenge is further complicated by overlaps with existing regimes. Healthcare technologies are already subject to the Medical Device Regulation (MDR) and the In Vitro Diagnostic Regulation (IVDR). The AI Act's requirements may duplicate or conflict with these frameworks, forcing developers and providers to navigate a complex regulatory landscape. Inspiration from the GDPR is evident in the Act's emphasis on data quality and transparency, but the interaction of these instruments can create uncertainty even as they strengthen patient protection (Brandt Brodersen et al., 2025).

The AI Act also has extraterritorial reach. Like the GDPR before it, its standards apply not only to providers based within the Union but also to non-EU actors whose systems are marketed in Europe. This "Brussels effect" ensures that the Act's influence will extend far beyond Europe's borders. By contrast, the United States relies on a fragmented, sectoral approach to AI oversight, with the Food and Drug Administration regulating some medical applications but leaving broader ethical and accountability issues unresolved. The divergence positions Europe as a global standard-setter for trustworthy AI, but leadership brings risks: excessive rigidity may drive innovation elsewhere and slow Europe's ability to benefit from breakthroughs (Linera & Meuwese, 2025, p. 1-4).

At its core, the regulation of healthcare AI is about more than compliance. It is about fundamental rights. Patients must be guaranteed privacy, fairness, and dignity in decisions that shape their health. The AI Act embeds these values into its requirements for transparency, oversight, and robustness, reflecting the EU's broader constitutional commitment to human rights. Yet whether these safeguards will work in practice remains uncertain. Healthcare systems already face resource shortages and staff pressures that complicate meaningful oversight (Solaiman, 2025, p. 7).

This article argues that the AI Act marks a turning point in the governance of healthcare AI. Its risk-based framework strengthens patient protections and affirms Europe's human rights commitments, but its automatic classification of healthcare AI as high-risk creates costs and complexities that disproportionately burden SMEs and smaller healthcare institutions. The Act's interplay with the MDR, IVDR, and GDPR further complicates compliance, and while it positions the EU as a global standard-setter, it also raises the risk of innovation flight.

To explore these issues in depth, the article is structured as follows. Chapter One explains the background and core principles of the AI Act, situating it within the EU's broader trajectory of digital regulation and examining the risk-classification system that underpins it. Chapter Two turns to the healthcare sector itself, analyzing the types of AI systems covered by the Act and the reasons they are automatically designated high-risk. Chapter Three considers the legal responsibilities imposed on different stakeholders, from developers and providers to healthcare institutions and professionals, including the requirements of transparency, data governance, and human oversight. Chapter Four examines the interaction between the AI Act and existing instruments such as the MDR and IVDR, assessing whether their combined application facilitates or complicates compliance. Chapter Five addresses the ethical and human rights implications of healthcare AI, including the risks of bias, opacity, and erosion of dignity. Chapter Six evaluates the practical challenges of compliance, particularly for SMEs and smaller providers, and proposes reforms such as regulatory sandboxes, tailored guidance, and funding support. Finally, it concludes by reflecting on the EU's role as a global standard-setter, highlighting both the opportunities and risks of its ambitious regulatory model.

2. BACKGROUND AND CORE PRINCIPLES OF THE AI ACT

The AI Act did not emerge in a vacuum but as the culmination of more than a decade of regulatory experimentation in the digital domain. The EU has long sought to establish itself as a global leader in digital rights and data governance, beginning most prominently with the GDPR, which set new standards for privacy and data protection worldwide. The AI Act represents the next step in this trajectory, expanding the Union's ambition beyond safeguarding personal data to governing the design and deployment of algorithmic technologies themselves. Its enactment in June 2024 was the product of years of negotiation among European institutions, national governments, industry stakeholders, and civil society organizations, all seeking to balance innovation with the protection of fundamental rights (Linera & Meuwese, 2025, p. 8).

At the core of the AI Act lies its risk-based approach, which classifies systems according to the potential harm they may cause. This reflects the EU's broader regulatory philosophy: not all technologies pose the same dangers, and law must calibrate obligations accordingly. At the top of the hierarchy are systems deemed to present "unacceptable risk," such as those involving

manipulative subliminal techniques or social scoring practices reminiscent of authoritarian regimes. These are banned outright under Article 5. Beneath them are "high-risk" systems, which include applications in critical infrastructure, law enforcement, migration control, and-most relevant here- healthcare. High-risk systems are permitted, but only under strict obligations designed to safeguard transparency, accountability, and human oversight. Other systems are classed as limited- or minimal-risk and face lighter requirements, though even minimal-risk systems may be subject to transparency or disclosure rules (Fraser & Bello y Villarino, 2023, p. 2-3).

Healthcare occupies a special position in this framework because AI systems in this sector are automatically classified as high-risk, regardless of their particular function or technical design. This classification reflects the Union's recognition that medical technologies directly affect life, health, and human dignity. It also demonstrates how the AI Act weaves technological governance together with constitutional values, embedding fairness, privacy, equality, and dignity into regulatory obligations. By treating healthcare as a domain demanding maximum precaution, the EU makes clear that innovation cannot come at the expense of patient safety or public trust (van Kolfschooten, 2024, p. 1-4).

The principles underpinning the AI Act extend beyond classification to include an explicit commitment to fundamental rights. The legislation draws on the EUCFR, the European Convention on Human Rights, and long-standing traditions of human dignity in European constitutional law. These rights are not abstract ideals but operational guideposts. Article 10 on data governance embodies equality and non-discrimination (Charter Art. 21) by mandating the use of relevant, representative, and bias-free datasets. Article 13 on transparency gives effect to privacy and autonomy rights (Charter Arts. 7–8) by requiring that individuals understand the decisions that affect them. Article 14 on human oversight reflects the conviction that machines should not replace human judgment in contexts involving life and dignity (Charter Arts. 1 and 3). Article 15, requiring accuracy, robustness, and cybersecurity, safeguards reliability and patient safety, linking directly to the Charter's guarantee of healthcare (Art. 35).

The AI Act also aspires to set global standards. Like the GDPR, it has extraterritorial scope, applying not only to providers established within the Union but also to those outside who place AI systems on the European market. This "Brussels effect" ensures that the Act's influence extends well beyond Europe's borders. In practice, this means that American, Asian, and other international firms seeking to access the EU market must align with its principles of trustworthy AI. By contrast, the United States has adopted a more fragmented, sectoral approach, relying on agencies such as the Food and Drug Administration to regulate certain applications while leaving broader accountability questions unresolved. Europe, in taking a

comprehensive and rights-based path, positions itself as the jurisdiction most willing to constitutionalise AI governance (Linera & Meuwese, 2025, p. 10; RAND Europe, 2024; Mello & Cohen, 2025).

The Act's development was also shaped by political dynamics. Public unease about AI was growing, fuelled by reports of biased algorithms, opaque decision-making, and dystopian scenarios in which machines eroded human autonomy. Civil society groups pushed for strong safeguards, while industry lobbied for flexibility to protect competitiveness. Member States brought divergent perspectives, some prioritising innovation, others emphasising ethics. The final text reflects compromise: ambitious in creating the first comprehensive framework of its kind, but pragmatic in providing for regulatory sandboxes and iterative guidance to ease compliance, particularly for SMEs (Brandt Brodersen et al., 2025).

One of the Act's most striking features is its insistence on explainability and documentation. High-risk systems must be accompanied by technical files (Art. 11) and event logs (Art. 12) that enable regulators and users to trace how systems function and how risks are managed. This reflects a European conviction that transparency is essential for accountability, but it also exposes the tension between law and technology: many machine learning systems, especially deep neural networks, resist simple explanation. The legal requirement for explainability thus becomes both a regulatory mandate and a technological challenge, pushing developers to design systems that are not only powerful but interpretable (Busch et al., 2024, p. 3-4).

Finally, the AI Act is closely tied to Europe's broader project of digital sovereignty. By setting harmonised rules for AI, the Union seeks not only to protect rights but also to assert autonomy in a field dominated by American and Chinese firms. For healthcare, this ambition is particularly important. By defining what counts as safe and trustworthy medical AI, Europe can shape the direction of global health innovation while ensuring that its constitutional values remain central to the technologies that affect patients' lives.(Busch et al., 2024, p. 3-4).

In sum, the background and principles of the AI Act reveal a convergence of legal, ethical, and political currents. The legislation builds on Europe's history as a regulator of digital rights, adopts a risk-based framework that calibrates obligations to potential harm, and embeds fundamental rights as its guiding star. Its extraterritorial scope gives it global influence, while its insistence on transparency, oversight, and robustness reflects a human-centered vision of technological governance. In healthcare, the automatic high-risk classification encapsulates the Union's ambition that patient safety and dignity must always remain paramount. These principles set the stage for the next chapter, which examines in greater detail the specific risks of healthcare AI and why the Act imposes its most stringent requirements in this sector.

3. AI IN HEALTHCARE AS AUTOMATICALLY HIGH-RISK

By classifying all healthcare applications of AI as "high-risk," the AI Act made a clear statement about the moral and legal significance of medicine. The regulation does not draw fine distinctions between diagnostic algorithms, administrative scheduling tools, or predictive models. Instead, it declares the sector as a whole subject to the strictest oversight, regardless of the specific function or technical sophistication of a given system. This broad stance reflects the Union's view that the stakes in healthcare are uniquely grave: errors can cost lives, compromise dignity, and erode trust (Vardas et al., 2025, p. 836-838).

The first justification for this blanket designation is the catastrophic impact of medical error. An algorithm that misclassifies an online shopping preference poses little harm; one that misclassifies a malignant tumour as benign could prove fatal. Policymakers concluded that AI in medicine cannot be regulated with the same tolerance as AI in finance, advertising, or logistics. A single erroneous output in a clinical setting may result in delayed treatment, misdiagnosis, or even death. The Union therefore embraced a precautionary stance: better to err on the side of overprotection than to risk under-regulation in life-and-death contexts (Mello & Cohen, 2025; Busch et al., 2024, p. 12).

A second concern lies in the opacity of advanced systems. Many machine learning models, particularly deep neural networks, can generate outputs that are accurate yet inscrutable. For doctors, this creates a dilemma. They are expected to justify clinical decisions to patients, colleagues, and courts. When a machine offers a recommendation without an intelligible explanation, accountability is disrupted. If a physician follows the output and harm occurs, reliance on the system may be faulted; if the physician disregards it and harm still results, refusal to follow may also be questioned. For patients, being told simply that "the algorithm decided" undermines informed consent. The AI Act addresses this problem by requiring technical documentation (Art. 11), record-keeping (Art. 12), and transparency obligations (Art. 13) that make these systems auditable and explainable to the extent technology allows.

Bias in datasets provides another strong rationale for automatic high-risk classification. Medical data are often fragmented, incomplete, and skewed by historical inequities. Women, minorities, and populations outside Europe and North America have long been underrepresented in clinical research. Algorithms trained on these data risk reproducing inequalities at scale. The widely cited case of pulse oximeters, which overestimated oxygen levels in patients with darker skin tones, illustrates how such bias can have deadly consequences. By treating healthcare AI as high-risk, the AI Act requires developers to apply rigorous data governance (Art. 10), document dataset sources (Art. 11), and actively mitigate discriminatory effects. These provisions give concrete form to the Charter's guarantees of equality (Art. 21) and the right to healthcare (Art. 35).

The Act also responds to the risk that professional judgment could be eroded. One of AI's great attractions in healthcare is its potential to relieve overburdened staff by triaging patients, supporting diagnoses, or generating treatment options. Yet efficiency can become dangerous if it undermines clinical autonomy. Doctors may feel pressured to follow algorithmic outputs simply because they appear authoritative or because institutions encourage reliance on them. Over time, this could reduce the role of human expertise in care. To resist such drift, Article 14 requires "meaningful human oversight" of high-risk AI, reaffirming that responsibility must remain with healthcare professionals.

The ethical dimension reinforces these justifications. Healthcare patients are often in vulnerable positions, reliant on compassion and trust. Introducing opaque, data-driven systems without adequate safeguards risks alienating patients and diminishing confidence in medical institutions. Guided by constitutional traditions of dignity, fairness, and equality, lawmakers considered such risks intolerable. The blanket high-risk classification is therefore more than a technical safeguard; it is a normative statement that innovation must remain subordinate to patient dignity (Solaiman, 2025, p. 18).

Critics argue that this approach is too blunt. Not all healthcare AI tools carry the same risk. A scheduling program for appointments should not face the same obligations as a system determining chemotherapy dosages. The AI Act, however, draws no distinction. Both are high-risk by default. This rigidity may impose unnecessary costs, slow the deployment of useful but low-stakes tools, and particularly burden SMEs, which often develop administrative or niche technologies (Brandt Brodersen et al., 2025).

Policymakers defend the broad approach on pragmatic grounds. Drawing lines in healthcare is notoriously difficult: an apparently low-risk tool today might, through integration, tomorrow influence diagnosis or treatment. By applying strict obligations universally, the Union reduces uncertainty and sends a clear message to citizens: healthcare AI will be subject to the highest standards. Politically, this choice builds trust. Citizens may not understand the inner workings of neural networks, but they understand a commitment to patient safety (RAND Europe, 2024).

The implications are extensive. Developers must implement risk management systems (Art. 9), maintain high-quality and bias-resistant data (Art. 10), prepare documentation (Art. 11), keep logs (Art. 12), ensure transparency (Art. 13), provide for human oversight (Art. 14), and guarantee accuracy, robustness, and cybersecurity (Art. 15). Providers and institutions must train staff, monitor system performance, and establish governance mechanisms. Compliance is costly and time-consuming, but its purpose is to ensure that patients receive care that is both technologically advanced and rights compliant.

Ultimately, the automatic high-risk designation is not simply a regulatory choice but a constitutional stance. It reflects the EU's determination to reconcile innovation with dignity, equality, and safety. Compliance costs may be high, particularly for smaller actors, but the alternative- an efficient yet unjust healthcare system governed by opaque and biased algorithms- would be fundamentally incompatible with European law and values.

4. LEGAL RESPONSIBILITIES OF STAKEHOLDERS

The AI Act does not merely classify healthcare applications as high-risk; it also redefines the legal responsibilities of those who design, distribute, and use them. This is one of the regulation's most transformative contributions. By creating detailed obligations for developers, providers, healthcare institutions, and professionals, the Act addresses accountability gaps that have long troubled debates about algorithmic medicine. Responsibility is no longer diffuse or ambiguous but distributed across the entire lifecycle of an AI system- from conception and training to deployment and post-market monitoring.

For developers and providers, the obligations are especially demanding. Before a high-risk AI system can be placed on the European market, providers must undergo conformity assessments (Arts. 43–51). These require evidence that the system complies with obligations relating to risk management (Art. 9), data governance (Art. 10), technical documentation (Art. 11), record-keeping (Art. 12), transparency (Art. 13), human oversight (Art. 14), and accuracy, robustness, and cybersecurity (Art. 15). Providers must document how data were collected, cleaned, and validated; maintain technical files describing the system's design and intended use; and implement quality management systems that monitor and mitigate risks throughout the lifecycle. These are not check-box requirements but ongoing commitments demanding substantial resources and expertise. Large corporations may absorb the costs, but for SMEs- the drivers of much health technology innovation- the burden can be overwhelming. The Act attempts to ease this through regulatory sandboxes (Arts. 53–57), yet the challenge remains formidable.

Transparency is a core duty. Under Article 13, developers cannot offer systems that generate outputs without explanation. They must supply documentation and plain-language instructions sufficient for healthcare professionals to understand limitations and for regulators to exercise oversight. Transparency thus becomes a shared obligation: developers must provide it, and users must engage with it. Without such reciprocity, explainability risks degenerating into mere formalism (Solaiman, 2025, p. 8-10).

Data governance is another pillar of responsibility. Article 10 requires that training datasets be relevant, representative, and sufficiently free from bias. This obligation recognises that biased data produce biased outcomes, and in healthcare those outcomes may be life-threatening.

Developers must actively identify and correct imbalances in their data. Achieving this is not simple: medical data are fragmented across institutions, shaped by historic inequities, and restricted by privacy rules. Meeting Article 10's standard requires major investment in infrastructure, coordination, and oversight, but the stakes justify the effort (Vardas et al., 2025, p. 837).

Healthcare institutions also carry heavy responsibilities. It is not enough for developers to deliver compliant systems; hospitals and clinics must ensure responsible use. Institutions must verify that systems have passed conformity assessments, train staff to understand both functionality and limitations, and establish monitoring processes to detect errors or anomalies. As "users" under Article 30, institutions must operate systems according to instructions, maintain records of use, and cooperate with post-market monitoring. They must also ensure that patients are informed about the role of AI in their care, aligning deployment with the principle of informed consent. For already resource-constrained institutions, these obligations are demanding but necessary to ensure that regulation translates into practice.

Individual healthcare professionals also retain responsibility. Article 14 emphasises "meaningful human oversight" in high-risk contexts, reaffirming that professionals cannot blindly defer to machines. Doctors, nurses, and other practitioners must exercise independent judgment, question outputs where appropriate, and ensure that patients are treated as individuals rather than datapoints. This preserves professional autonomy and ethical responsibility but also raises difficult liability questions. If harm results from following an AI recommendation, is the doctor, the institution, or the developer liable? The AI Act does not answer fully, but it establishes the framework within which courts and regulators will determine responsibility.(Solaiman, 2025)

Regulators and notified bodies likewise play an essential role. Independent bodies must conduct conformity assessments, while market surveillance authorities must enforce compliance. Their effectiveness will depend on resources and expertise. Recognising this, the AI Act created the European AI Office (Art. 56) to coordinate supervision across Member States and harmonise standards. Without sufficient capacity, enforcement risks becoming symbolic rather than substantive.

The distribution of obligations reflects a broader regulatory philosophy: AI governance in healthcare must be polycentric. Developers must ensure safety by design, institutions must guarantee responsible deployment, professionals must preserve ethical application, and regulators must enforce oversight. Patients themselves indirectly shape accountability by demanding transparency and fairness from the systems that affect them. The Act weaves these responsibilities together into a fabric of governance intended to preserve trust while enabling innovation (Goisauf, 2025, p. 5).

The weight of these obligations, however, is considerable. SMEs face disproportionate difficulties. A small company designing an AI tool for patient monitoring may lack the resources to audit datasets, establish quality management systems, and prepare extensive documentation. Without targeted support, such firms risk exclusion, leading to concentration of innovation among large corporations. This raises a policy dilemma: how to uphold high standards without creating insurmountable barriers for smaller actors. Regulatory sandboxes offer one response, but they must be complemented by funding support and clear, accessible guidance (Brandt Brodersen et al., 2025; RAND Europe, 2024).

In the end, the AI Act's distribution of responsibilities reflects Europe's normative choice to embed accountability into every stage of the AI lifecycle. Trust cannot be outsourced to machines; it must be maintained through human involvement, transparency, and fairness. The obligations of Articles 9 through 15, together with those of Articles 29 and 30, are not mere technicalities but expressions of constitutional values. By turning abstract rights into concrete duties, the AI Act ensures that artificial intelligence in healthcare remains consistent with the Union's deepest commitments to dignity and justice.

5. INTERPLAY WITH MDR AND IVDR

The AI Act does not operate in isolation but alongside pre-existing frameworks that already govern medical technologies. Two of the most significant are the MDR and the IVDR, both of which impose strict requirements on safety, performance, and market access for healthcare products. The AI Act overlays these frameworks with additional obligations that address the distinct risks posed by algorithmic systems: opacity, bias, lack of human oversight, and cybersecurity. This interaction creates a dense regulatory landscape that is both comprehensive and, at times, difficult to navigate (van Kolfschooten, 2024).

Under Annex III of the AI Act, any AI system that qualifies as a medical device or in vitro diagnostic device under the MDR or IVDR is automatically classified as high-risk. As a result, providers of such technologies must comply not only with device-specific rules under the MDR or IVDR but also with the AI-specific requirements of Articles 9 through 15. For example, an AI-driven diagnostic tool must satisfy the MDR's safety and performance requirements while also demonstrating compliance with the AI Act's obligations on risk management, data governance, documentation, transparency, human oversight, and robustness. The outcome is a dual compliance burden: one regime ensures physical safety and technical reliability, the other ensures algorithmic accountability and rights protection (Busch et al., 2024; van Kolfschooten, 2024).

The risk of duplication is significant. MDR and IVDR already require conformity assessments, quality management systems, and post-market surveillance. The AI Act introduces parallel obligations that may overlap or conflict. Article 73 attempts to mitigate this challenge by requiring cooperation between AI and medical device regulators. In theory, such coordination should prevent double assessments and harmonise standards. In practice, however, aligning procedures across Member States and across different regulatory authorities will be complex. Without careful implementation, developers may face duplicative audits, extended approval timelines, and inconsistent interpretations of their obligations (Mello & Cohen, 2025).

The European AI Office, established under Article 56, is intended to play a central role in this process. Acting as a supervisory and coordinating body, it is tasked with ensuring that AI-specific oversight is integrated effectively with sectoral regimes such as MDR and IVDR. Its guidance and dialogue with national authorities will be critical to preventing regulatory fragmentation. If the Office succeeds, Europe could develop a model where different frameworks complement one another to produce both safe and trustworthy healthcare technologies. If it fails, Europe risks building a system where innovators and providers face confusion, inefficiency, and delay (Kutscher, 2025).

For developers, this interplay means preparing compliance materials that address both sets of obligations. Technical documentation under the MDR must be supplemented with algorithmic documentation under Article 11 of the AI Act. Post-market surveillance under the MDR must be paired with logging and monitoring under Articles 12 and 61 of the AI Act. Healthcare institutions, too, must ensure that tools they adopt meet dual compliance requirements. For large corporations, this may be achievable through specialised compliance teams. For SMEs, however, the complexity and cost of satisfying overlapping rules may create insurmountable barriers, stifling innovation (Brandt Brodersen et al., 2025).

Yet this dual framework is also a deliberate normative choice. The EU has signalled that medical technologies must not only function safely but also respect fairness, transparency, and human dignity. By layering the AI Act onto MDR and IVDR, the Union affirms that technical safety is necessary but not sufficient; algorithms that affect human health must also be accountable, interpretable, and rights compliant. The price of duplication is therefore the cost of embedding constitutional values directly into the governance of medical technologies (Solaiman, 2025).

In sum, the interplay between the AI Act and the MDR/IVDR reflects both the ambition and the difficulty of Europe's approach. If regulatory coordination under Article 73 and the AI Office functions effectively, Europe could produce a coherent "gold standard" for trustworthy medical AI. If coordination falters, however, the result may be duplication, inefficiency, and barriers to innovation, particularly for SMEs. The outcome will depend less on the words of the legislation than on the capacity of institutions to harmonise its application across the Union (Vardas et al., 2025).

6. ETHICAL AND HUMAN RIGHTS DIMENSIONS

The regulation of artificial intelligence in healthcare is not only a matter of technical compliance but one of fundamental rights and ethical principles. Unlike in other sectors, where AI may serve as an efficiency tool, in healthcare it implicates dignity, autonomy, equality, and the trust between patient and physician. When algorithms misclassify tumors, produce biased recommendations, or conceal their logic, the harm is not limited to technical error—it strikes at the very foundation of human rights and undermines confidence in medical institutions. The AI Act recognizes this reality by embedding rights and ethics into its legal structure, making healthcare AI governance a constitutional project as much as a regulatory one (Solaiman, 2025; Goisauf, 2025).

The Act draws its most categorical line in Article 5, which prohibits AI practices deemed to present "unacceptable risk." These include systems that manipulate human behavior through subliminal techniques, exploit vulnerabilities, or engage in social scoring. In healthcare, where patients often confront illness in states of dependence and vulnerability, such prohibitions take on heightened importance. If AI were deployed to steer patients toward decisions that served corporate interests at the expense of their welfare, such systems would violate not only statutory norms but also the Charter's guarantee of human dignity (Art. 1) and the integrity of the person (Art. 3). Article 5 therefore functions as a constitutional safeguard, establishing that certain practices are never permissible regardless of efficiency or potential benefit (Charter of Fundamental Rights of the European Union, 2012).

Beyond prohibitions, the AI Act sets rigorous positive obligations designed to translate rights into operational duties. Article 10 requires that high-risk AI systems be trained, validated, and tested on datasets that are relevant, representative, and sufficiently free from bias. This obligation reflects the Charter's guarantees of equality and non-discrimination (Art. 21) and the right to healthcare (Art. 35). The lesson of biased datasets—such as the well-documented disparities in medical devices like pulse oximeters—underscores that data inequities become health inequities when encoded into AI. By mandating high standards of data governance, Article 10 turns the principle of fairness into a legal requirement (Vardas et al., 2025).

Article 11 on technical documentation and Article 12 on record-keeping reinforce accountability by ensuring that AI systems are not deployed as opaque black boxes. Developers must produce technical files and maintain logs that allow regulators and users to trace how systems function and how risks are managed. These duties reflect the Charter's rights to privacy (Arts. 7–8) and to an effective remedy (Art. 47), because rights cannot be exercised if affected individuals cannot understand or challenge the systems that influence their care (Mello & Cohen, 2025).

Transparency obligations under Article 13 strengthen this connection. Patients and professionals must be informed about the capabilities, limitations, and intended uses of AI systems. In practice, this means that doctors should not receive algorithmic recommendations without clear explanations of their basis and boundaries, and patients must not be treated as passive subjects of machine logic. Transparency thus functions as a bridge between technical complexity and human agency, enabling meaningful informed consent and preserving patient autonomy (Busch et al., 2024).

Article 14 further requires meaningful human oversight of high-risk systems. Healthcare professionals must retain the ability to monitor, interpret, and override AI outputs where necessary. This provision protects not only professional autonomy but also the dignity of patients, who are entitled to human judgment in decisions affecting their lives. Oversight obligations directly address the risk of automation bias, the tendency for humans to defer uncritically to machine-generated outputs. In healthcare, such deference could have fatal consequences. Article 14 therefore ensures that accountability remains human-centered, consistent with the Charter's values of dignity and integrity (Solaiman, 2025).

Completing this framework, Article 15 requires that high-risk AI systems achieve high standards of accuracy, robustness, and cybersecurity. In healthcare, where errors may be fatal, robustness is not merely a technical aspiration but a human rights imperative. Systems must be resilient against malfunction and secure against attack, for breaches in accuracy or cybersecurity can directly endanger patient lives. By embedding these obligations, the AI Act affirms that the reliability of healthcare AI is a matter of constitutional significance, tied to the Charter's protection of life and health (van Kolfschooten, 2024).

Taken together, these provisions form an interlocking set of safeguards: prohibitions on unacceptable practices (Art. 5), fair and representative data (Art. 10), documentation and record-keeping (Arts. 11–12), transparency (Art. 13), human oversight (Art. 14), and robustness and cybersecurity (Art. 15). Each resonates with the Charter of Fundamental Rights: dignity (Art. 1), integrity of the person (Art. 3), privacy (Arts. 7–8), equality (Art. 21), healthcare (Art. 35), and effective remedies (Art. 47). The result is a framework in which healthcare AI is governed not only as a technical system but as a site where constitutional rights are directly at stake (Fraser & Bello y Villarino, 2023).

The ethical significance of this framework lies in its insistence that trust cannot be outsourced to machines. Patients must be assured that AI systems serve their interests, not those of corporations or governments; professionals must know that their expertise is reinforced, not undermined, by technological tools; and regulators must demonstrate that their oversight is substantive rather than symbolic. Legitimacy in healthcare depends on trust as much as on accuracy. By embedding rights into binding obligations, the AI Act attempts to ensure that trust in medicine is preserved even as algorithms play an expanding role (Goisauf, 2025).

In sum, the ethical and human rights dimensions of the AI Act show that the Union's approach to healthcare AI is about more than managing risk. It is about embedding dignity, equality, autonomy, and fairness into the technological infrastructures of modern medicine. Compliance with Articles 5, 10, 11, 13, 14, and 15 may be costly and demanding, but the alternative—an algorithmic healthcare system that is efficient yet unjust—would undermine the very values on which European healthcare rests. By constitutionalizing obligations, the AI Act charts a path for medical AI that is not only effective but also profoundly human (Solaiman, 2025; Brandt Brodersen et al., 2025).

7. PRACTICAL CHALLENGES AND REFORMS

The AI Act introduces a new era of regulatory complexity for healthcare innovation. While its objectives of safety, transparency, and accountability are widely supported, the practical implementation of these principles presents formidable challenges. Compliance obligations are extensive, costly, and ongoing. Duplication with existing regimes such as the MDR and IVDR risks inefficiency, and enforcement capacity across Member States remains uncertain (Mello & Cohen, 2025; Brandt Brodersen et al., 2025). Yet these difficulties do not negate the Act's importance—they highlight the need for pragmatic governance capable of balancing rights protection with innovation (Solaiman, 2025).

Compliance under Articles 9 through 15 demands substantial technical, legal, and administrative expertise. Developers must maintain risk management systems, ensure data quality, prepare documentation, monitor post-market performance, and guarantee human oversight. For large corporations, these duties may be absorbed through compliance departments and regulatory teams. For SMEs, however, they represent potentially existential burdens (RAND Europe, 2024). Many small innovators lack the staff, funding, or infrastructure to meet the AI Act's detailed technical requirements, raising the risk that only well-resourced firms can survive in this environment (Busch et al., 2024).

To address these asymmetries, the AI Act provides for regulatory sandboxes (Articles 53–57), designed to allow supervised experimentation in a controlled setting. These mechanisms permit developers—particularly SMEs—to test systems under regulatory supervision, gaining feedback without incurring full compliance costs prematurely. Sandboxes can build mutual understanding between innovators and regulators, but their effectiveness will depend on accessibility and consistent application across Member States (Brandt Brodersen et al., 2025).

Funding and guidance will also be crucial. The European Commission and Member States may need to establish dedicated funds or incentives to help smaller firms comply with the Act's complex requirements. Clear, harmonised guidance documents and sector-specific templates could simplify compliance without lowering standards (Kutscher, 2025). Public-private partnerships may further ease burdens by pooling expertise and infrastructure for risk management and testing (Linera & Meuwese, 2025).

Another challenge lies in enforcement. Effective supervision requires well-trained regulators capable of understanding advanced machine learning systems. Member States will need to invest in technical expertise within notified bodies and market surveillance authorities (Vardas et al., 2025). The European AI Office, created under Article 56, will play a central role in coordinating oversight and ensuring that national authorities apply consistent standards. Without adequate resourcing, however, enforcement risks becoming symbolic rather than substantive (van Kolfschooten, 2024).

Finally, practical implementation must align with the AI Act's ethical foundations. The regulation is not solely about technical control but about safeguarding dignity, fairness, and human rights in medical innovation. To preserve legitimacy, compliance processes should remain transparent, proportionate, and supportive of responsible innovation (Goisauf, 2025). The goal should be not to deter progress but to ensure that progress serves people (Solaiman, 2025).

8. CONCLUSION

The AI Act marks a turning point in the governance of technology. For the first time, a jurisdiction has not only sought to regulate artificial intelligence for safety but also to embed ethical and constitutional values into binding law. This ambition is most visible in healthcare, where AI systems are automatically categorised as high-risk under Annex III. In doing so, the Union made clear that the protection of patients- over efficiency, cost, or innovation speed- is its foremost priority.

The preceding chapters have shown how the AI Act operates on multiple levels. At its foundation, Articles 5 and 6 create the architecture of risk classification, prohibiting practices deemed to present unacceptable risks and subjecting high-risk systems to strict oversight. In healthcare, the automatic high-risk classification affirms that errors in this field carry consequences too serious to leave to voluntary standards.

At the operational level, Articles 9 through 15 create a demanding compliance ecosystem. Systems must be subject to risk management (Art. 9), built on fair and representative data (Art. 10), accompanied by technical documentation (Art. 11) and logs (Art. 12), transparent in their outputs (Art. 13), supervised by human judgment (Art. 14), and robust, accurate, and secure (Art. 15). These obligations transform fairness, transparency, and accountability from abstract principles into enforceable duties.

Responsibility is also distributed across the ecosystem. Articles 29 and 30 impose obligations not only on providers but also on users, including healthcare institutions and professionals. Article 61 extends oversight through continuous post-market monitoring, while Articles 64 to 69 empower regulators to impose penalties, including significant fines, for non-compliance.

Article 56 establishes the European AI Office to coordinate implementation, and Article 73 requires cooperation with existing regimes such as the MDR and IVDR. Together, these provisions build a framework of shared responsibility, in which accountability is not concentrated in a single actor but distributed across developers, institutions, professionals, and regulators.

Yet the strength of the Act produces undeniable strain. Compliance with Articles 9–15 and conformity assessments under Articles 43–51 requires substantial expertise and resources. For large multinational corporations, these demands may be manageable. For SMEs and smaller healthcare providers, they may be prohibitive. Regulatory sandboxes under Articles 53–57, along with tailored guidance and targeted funding, are crucial to preventing the exclusion of these innovators. Without such support, the AI Act risks consolidating the market around large incumbents and narrowing the diversity of available healthcare technologies.

Globally, the AI Act projects Europe as a regulatory leader. Its extraterritorial reach ensures that any provider seeking access to the Union's market must comply with its rules. This "Brussels effect," familiar from the GDPR, positions Europe as a global standard-setter for trustworthy AI. Leadership, however, brings risks: if implementation proves too rigid or costly, innovation may gravitate toward more permissive jurisdictions. Europe's challenge is to balance ambition with pragmatism, ensuring that its standards protect patients without deterring beneficial advances.

Ultimately, the AI Act is about more than compliance. It is about constitutional commitments: to dignity (Charter Art. 1), to integrity of the person (Art. 3), to privacy and data protection (Arts. 7–8), to equality (Art. 21), to healthcare (Art. 35), and to effective remedies (Art. 47). By embedding these rights into binding obligations, the Act ensures that healthcare AI is governed not only by technical standards but by values that define Europe's legal order.

If implemented wisely- through strong enforcement, accessible guidance, robust regulatory capacity, and support for smaller actors- the AI Act could establish a gold standard for trustworthy medical AI. It would demonstrate that innovation can be aligned with dignity, transparency, and fairness, and that technology need not come at the cost of human rights. If implemented poorly, however, the Act risks creating barriers, stifling innovation, and undermining its own legitimacy.

The stakes are immense. The future of medicine will be inseparable from the future of artificial intelligence. By insisting that AI serve humanity rather than the reverse, the AI Act offers a vision of a healthcare system that is both technologically advanced and profoundly human. It is not the final word on AI governance, but it is a decisive first step—an opening chapter in the effort to reconcile innovation with justice, efficiency with dignity, and law with rapid technological change.

REFERENCES

AI Act Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on artificial intelligence (AI Act).

Brodersen, J. B., Caggiano, I. A., Kringen, P., Madai, V. I., Osika, W., Sartor, G., Svensson, E., Westerlund, M., & Zicari, R. V. (2025). Getting Ready for the EU AI Act in Healthcare: A call for Sustainable AI Development and Deployment. arXiv:2505.07875.

Busch, F., et al. (2024). Navigating the European Union Artificial Intelligence Act for healthcare. *npj Digital Medicine, 7*, 210.

Charter of Fundamental Rights of the European Union. (2012). 2012 O.J. C 326/391.

Fraser, H., & Bello y Villarino, J. M. (2023). Acceptable risks in Europe's proposed AI Act. *Preprint*.

General Data Protection Regulation (GDPR). Regulation (EU) 2016/679.

Goisauf, M. (2025). Trust, trustworthiness, and the future of medical AI. *JMIR*.

In Vitro Diagnostic Regulation (IVDR). Regulation (EU) 2017/746.

Kutscher, S. (2025). The EU AI Act: Law of unintended consequences?

Medical Device Regulation (MDR). Regulation (EU) 2017/745.

Mello, M. M., & Cohen, I. G. (2025). Regulation of health and health care artificial intelligence. *JAMA*.

Presno Linera, M. Á. (2025). Regulating AI from Europe: A joint analysis of the AI Act.

Price, N. W. II. (2024). Liability for use of artificial intelligence in medicine. *Michigan Law School Faculty Publications*.

RAND Europe. (2024). *Risk-based AI regulation: A primer on the Artificial Intelligence Act*.

Solaiman, B. (2025). Regulating algorithmic care in the European Union: Evolving challenges under the AI Act. *Medical Law Review, 33*(1).

Vardas, E. P., et al. (2025). Medicine, healthcare and the AI Act: Gaps, challenges and consequences. *European Heart Journal – Digital Health*.

van Kolfschooten, H. (2024). The EU Artificial Intelligence Act: Implications for the healthcare sector. *Health Policy*.

ABOUT AUTHOR

Asst. Dr. Samet TATAR | Ankara Yıldırım Beyazıt University | samettatar[at]gmail.com | ORCID: 0000-0002-3553-1683

After graduating from the Faculty of Law at Istanbul University, he completed his master's and doctoral studies in the United States. He currently serves as the Head of the Department of Information Technology Law at Ankara Yıldırım Beyazıt University (AYBU). Additionally, he is the Director of the Institute of Information Technology at Tashkent State University and also serves as the Director of the Turkish-Uzbek Research Center jointly established by AYBU and Tashkent State University of Law. His academic interests include internet law, artificial intelligence law, and disinformation law.