THE FOUNDATIONAL DOCUMENTS SERIES · Vol. I ¦ No. I ¦ November 2024

WHITE PAPER

Issued by the Founding Council of the International Society of Medical AI (ISMAI)

Doctors Must Control Medical AI

not be controlled by it

ISMAI

International Society of Medical AI

Machinam Medicorum

 Published by ISMAI · www.ismai.org

Executive Summary

The rise of artificial intelligence (AI) in medicine is not a distant possibility. It is here, and it is sweeping and transformative. It has the capacity to illuminate diagnosis, streamline decision-making, and extend the reach of healthcare. Its promise, though, if unguided by human hands, becomes perilous. As AI assumes an expanding role at the bedside, the defining question of our era is clear: who holds the reins?

The International Society of Medical AI (ISMAI) was founded to ensure that physicians remain in control. This white paper does not position clinicians in opposition to technology, nor does it cast innovation as a threat to tradition. Artificial intelligence and medical ethics must be braided together, not torn apart. AI is not the adversary of clinical practice. Rather, it can be its ally, but only if it is governed, not by the loudest voices in the marketplace, but by those who are bound to the patient by oath, training, and accountability.

ISMAI offers a bridge between two essential worlds: that of the engineer and that of the clinician. It rejects the false binary of acceleration versus caution and proposes instead a shared architecture of trust. Hospitals become AI-ready not by merely installing software but by fostering a culture of transparency, shared governance, and informed consent in which algorithms are not deployed in isolation but rather are co-designed with those who will use them to protect human lives.

This vision rests on several pillars. First, clinical leadership must anchor the development of AI tools to ensure they meet standards of fairness, safety, and clinical relevance before they are used. Second, implementation must be matched by institutional readiness: governance structures, legal clarity, and cultural acceptance are foundational, not optional. Third, adaptation must be continuous. AI evolves; so too must our systems of validation and oversight. Continuous monitoring, open feedback loops, and ethical vigilance are essential if these technologies are to remain worthy of the trust patients place in the medical profession.

ISMAI convenes physicians, engineers, ethicists, regulators, and patients in a shared commitment to ensure that AI in medicine reflects not only what is possible but also what is right. Through certification programmes, position papers, and interdisciplinary education, including the development of certificate, diploma, and MSc pathways, ISMAI empowers professionals to become stewards of this transition rather than bystanders.

The stakes are immense, but so is the opportunity. We need not choose between human insight and computational power; when guided wisely, they can complement each other. ISMAI invites all who share its values to join a project larger than any algorithm: the safeguarding of care. In the space between innovation and tradition, between data and dignity, we must not allow a chasm to form. We must build a bridge.

I. The Dawn of AI in Healthcare

Artificial intelligence in healthcare has arrived with such speed and complexity that many find themselves unprepared for its profound implications. In one moment, clinicians are promised unprecedented diagnostic power and workflow efficiencies; in the next, they confront black-box algorithms that threaten to supplant rather than support their clinical judgments. In this atmosphere of both promise and peril, it has become clear that the medical profession can no longer afford to be passive. The future of healthcare *will* be shaped by AI, whether doctors guide this transformation or not, simply acquiescing to decisions made largely outside the profession's sphere of influence. To allow others, perhaps lacking the ethical imperatives and accountability that define medical practice, to determine how AI is used at the bedside is to cede control over patient welfare, safety, and trust. If physicians do not stand at the helm of this endeavour to guide it, the swift currents of technology may pull healthcare into places it was never meant to go. To the detriment of all.

ISMAI was formed to respond to this pivotal moment. Its vision is that AI, if harnessed responsibly, can assist clinicians and patients, reducing errors, streamlining care, and democratising access to advanced diagnostics. However, these benefits are sustainable only if AI is subject to the ethical, legal, and clinical rigor that has long guided medicine. Without robust validation and governance, AI will endanger patient safety and undermine public confidence in medicine. Human lives and dignity are at stake.

The medical community should not condemn innovation, of course. It must champion a balanced pathway forward. ISMAI is an international society where engineers, clinicians, ethicists, policymakers, and patient advocates come together under a single guiding principle: patient well-being is primary. It seeks to serve as a neutral liaison between AI developers and healthcare providers. By articulating rigorous standards, providing certification programs, and shaping clear governance structures, ISMAI aims to ensure that technological advances do not eclipse the deeply human core of clinical practice.

Physicians have a duty of care shaped by centuries of cumulative wisdom, and that duty must extend into the era of machine learning and data-driven decision support. If AI systems fail to reflect the core medical virtues of compassion, equity, and accountability, no technical sophistication can redeem their misuse.

ISMAI calls on every clinician, researcher, and concerned citizen to engage in this mission. Its purpose is not to stifle AI's evolution but to guide it towards a future in which doctors work hand in hand with intelligent tools, retaining ultimate accountability for their patients' outcomes. Only through principled leadership can we ensure that this rapidly unfolding wave of technology enriches healthcare rather than distorts it, and that tomorrow's triumphs remain firmly anchored in the ethics that have long defined the healing arts.

II. The Clinician's Voice Must Prevail

Clinicians stand at the intersection of science and empathy, entrusted by society to safeguard human health through reasoned judgment and moral responsibility. This duty cannot be delegated to opaque systems or corporate strategies devoid of direct accountability for patient welfare. When medical decisions are made by entities that do not adhere to the exacting standards of clinical practice, the implications for patient safety and

public trust are profound. Ethically, medicine is rooted in a commitment to, first, *do no harm*, and to respect each patient's autonomy and dignity. Artificial intelligence, however advanced, must be governed by those who live by these values. If algorithms are allowed to direct care without meaningful clinical oversight, the moral chain of accountability will break, leaving no clear guardian of patient interests.

The legal landscape likewise reflects a world unprepared for the complexities introduced by self-learning AI tools. Existing liability doctrines, fashioned for human decisions and physical devices, offer little clarity for cases in which clinicians have relied on proprietary algorithms. Without careful regulation, patients harmed by machine-generated errors could find themselves with no recourse and physicians will be unsure about the degree of responsibility they bear. Such uncertainty threatens the very trust upon which clinical interactions depend. If doctors fail to fully understand, or hesitate to embrace, the tools they are asked to use, or if patients sense that their caregivers are no longer fully in control of diagnostic and therapeutic processes, medicine will enter a new era of suspicion and hesitancy. Public trust, once eroded, is not easily reclaimed.

From a clinical perspective, unaccountable AI governance introduces unacceptable risks into diagnosis and treatment recommendations. AI systems have been shown to replicate biases in their data sources and thereby reinforce existing disparities in healthcare outcomes. Further, algorithms have delivered hazardous advice, necessitating human intervention to avert potential harm. Physicians must do more than serve as a final checkpoint in a chain of automated decisions. Good medical practice requires empathy, nuance, and contextual insight, and clinicians must proactively lead AI governance to ensure that the technology enhances rather than displaces these essential human dimensions.

People willingly entrust their lives to clinicians because they know that medical professionals are licensed, extensively trained, and bound by ethical codes that uphold the patient's well-being above all else. Allowing nonclinical entities to dominate medical AI governance would undermine the bond between patient and physician by introducing intermediaries with vastly different incentives.

In response to these converging pressures, ISMAI contends that clinical governance over AI is neither optional nor a mere best practice: It is an imperative grounded in the ethical, legal, and clinical realities that define modern healthcare. By bringing together physicians, AI developers, and ethicists under a common framework of accountability, ISMAI serves as a neutral liaison that seeks neither to halt innovation nor to uncritically yield to its momentum. Instead, it champions a vision of healthcare guided by the responsibilities, insights, and compassion that have shaped medical practice for centuries. Through continued dialogue, certification programs, and collaborative research, ISMAI aims to ensure that AI serves as a trusted ally in patient care.

III. Protect Data Privacy

In the history of medicine, each epoch-making innovation has required more than technical prowess or institutional ambition: It has relied on people who could interpret unfamiliar tools in the light of enduring clinical values and responsibilities. The marriage of artificial intelligence with healthcare marks the latest and perhaps most challenging of these moments. While AI developers gravitate towards speed, scale, and the transformation of data into predictive insights, healthcare practitioners are responsible for real-world outcomes, professional accountability, and

patient trust. In the gap between these two worlds, vital questions arise: How can an algorithm's capabilities be harnessed without eroding the clinician's moral and legal obligations? Who ensures that AI-driven decisions are transparent to those whose lives they affect? And how do we keep pace with the rapid evolution of machine learning while upholding the essential ethical codes that shape medical practice?

ISMAI exists precisely to navigate that intersection. It seeks to replace any latent antagonism between technological progress and clinical caution with a framework in which these aims illuminate and strengthen one another. Collaboration is the key, but it must be purposeful and guided by the principle that human well-being and dignity stand at the core of healthcare. When clinicians and AI developers work in isolation, they risk perpetuating misalignments: sophisticated algorithms may go unchecked by real-world testing or moral oversight, and clinical institutions may, fearing the unknown, dismiss potentially lifesaving innovations. ISMAI aims to establish channels for regular dialogue and transparent review that will foster a culture of open inquiry and mutual respect. It encourages developers to incorporate input from frontline medical staff at each stage (from data selection and model training to pilot deployment) and encourages hospitals and clinicians to adopt an agile mindset that can adapt to fast-paced AI research.

Central to ISMAI's approach is the idea of AI-ready hospitals. Becoming AI-ready is not merely a matter of installing new software or recruiting a data science team. It entails the creation of institutional structures and processes that uphold safety, accountability, and equity at every point of AI usage. AI-ready hospitals adopt transparent governance models that define roles and responsibilities for data custodians, clinicians, and AI specialists. These institutions endorse continuous performance monitoring, scrutinizing algorithms in light of evolving patient demographics and clinical guidelines. They also cultivate a culture of patient engagement in which patients are informed about the role AI might play in their diagnosis or treatment and retain the agency to seek human intervention whenever they feel uncertain.

This framework benefits developers as much as it does healthcare providers. Hospitals with robust oversight procedures and dedicated AI committees offer vital feedback loops for AI engineers, enabling rapid identification of biases, anomalies, or diminishing performance over time. Such environments promote iterative improvement rather than passive deployment, helping technology companies refine their products to meet clinical needs more precisely. Instead of responding defensively to scrutiny, AI developers can come to view these checks as indispensable to the building of tools worthy of widespread trust. Meanwhile, clinical leaders gain the reassurance that any AI tool involved in patient care has been subjected to rigorous testing, aligned with professional best practices, and continually evaluated for safety and efficacy.

The transition towards this collaborative future will not be without complexities. Some hospitals may be reluctant to invest in the necessary technological infrastructure, fearing that AI's potential is outweighed by privacy concerns or ethical uncertainties. Others might rush headlong into adopting new tools without building the supportive governance structures needed to maintain patient safety. ISMAI's role is to guide institutions away from both extremes. This task requires ongoing education of clinicians to ensure they understand the nature of algorithmic outputs and can interpret them effectively. It also calls for humility on the part of developers, who must recognise that success in healthcare demands not merely technical excellence but an abiding respect for human values.

In bridging these two worlds, ISMAI aspires to convert potential discord to constructive

synergy. The ultimate beneficiaries of this partnership will be the people whom medicine exists to serve: patients, whose lives and well-being depend on a union of scientific progress and compassionate care.

IV. Towards a New Framework of Implementation

ISMAI's vision arises from a fundamental belief that technology should serve as a partner rather than a prescriber. AI can strengthen diagnostic accuracy, streamline workflows, and enrich clinical decision-making, but these benefits materialise only when the frameworks guiding AI are anchored in medical ethics and professional accountability. Our era has witnessed both triumphant advances and cautionary setbacks in digital health, each underscoring the importance of structures that keep patients' well-being at the fore. AI's potential can be realised safely and equitably only if it is implemented under standards that accommodate rigorous validation and place clinicians at the helm.

ISMAI proposes that there must be clear criteria for selecting and validating AI tools before they enter clinical practice. Although technology vendors play an indispensable role in designing algorithms, the healthcare community remains the ultimate authority on patient care. An independent process for certifying hospitals and evaluating AI solutions is therefore crucial. By developing shared benchmarks, ISMAI aspires to give healthcare institutions a stable foundation from which to judge not only an algorithm's theoretical promise but also its real-world performance across different patient populations and care settings. One of ISMAI's core functions will be to define best practices for testing these tools in controlled environments, ensuring they meet acceptable thresholds of accuracy and fairness before broader deployment.

Cultural acceptance of AI is key, given that healthcare has long depended on the human bond between clinician and patient. AI systems might streamline administrative tasks and illuminate hidden patterns in data, but they must do so in a way that reinforces, rather than diminishes, the empathy and trust essential to medicine's deeper ethos. Equally imperative is legal preparedness, given that the existing liability landscape often appears ill-suited to addressing the subtleties of algorithmic recommendations. If clinicians are to remain the final arbiters of patient management, the law must adapt to allow for oversight that neither stifles innovation nor exposes practitioners and patients to undue risk. ISMAI will work with authorities and professional bodies to reconcile these concerns and to support a legal environment in which AI can flourish responsibly.

Our framework also accounts for the rapid pace at which digital health evolves. AI systems are seldom static; as they learn from new data or receive updates, they can change in ways unseen in traditional medical devices. ISMAI advocates a structure in which hospitals are equipped to conduct continuous monitoring to ensure that algorithms maintain their intended performance. This ongoing vigilance requires collaborative pathways between software developers, clinical teams, and regulatory agencies that are guided by conscientious governance committees at the institutional level. Through continual review, bias testing, and performance audits, healthcare systems will cultivate a culture that treats AI not as an unquestionable oracle but as a dynamic tool requiring perpetual calibration.

ISMAI's certification programs for hospitals, coupled with a suite of standardised guidelines for AI tools, aim to fill gaps in current policy while allowing room for further innovation. By taking a neutral stance—neither beholden to corporate imperatives nor dismissive of technological progress—ISMAI seeks to unify physicians, AI experts, and policy

leaders. We invite all healthcare stakeholders to engage in this model of shared governance, for it ensures that ethical values, patient safety, and medical expertise remain the bedrock upon which AI's promise is realised.

V. Who We Are and What We Offer

ISMAI was born of a conviction that medicine must preserve its commitment to healing, even in an era increasingly shaped by algorithms and data-driven innovation. Our founding core comprises physicians and other healthcare professionals who have firsthand knowledge of clinical workflows, patient vulnerabilities, and the complex ethical landscapes that underlie medical practice. Alongside these clinical experts, we have welcomed engineers and data scientists who bring depth in algorithmic research, ethicists who anchor new technologies in moral principles, and administrators who understand the realities of hospital governance. This inclusive framework allows us to address the challenges of artificial intelligence from every angle, ensuring that implementation decisions do not neglect any of the domains (clinical, technical, legal, and moral) that bear on patient outcomes.

At the heart of ISMAI's purpose is the creation of a trusted, neutral forum where diverse stakeholders can collaborate to shape the future of healthcare AI. By forming strategic alliances with hospitals and other care settings, we aim to act as both guide and partner, developing standards that equip healthcare leaders to make thoughtful, evidence-based choices about how AI is introduced to their wards, clinics, and operating rooms. These relationships serve to align technology firms with real-world clinical priorities and to harness innovation in ways that support, rather than supplant, the role of dedicated medical professionals.

One of our principal contributions to this evolving landscape is the design and dissemination of guidance, from scholarly position papers to hands-on toolkits, each addressing a different facet of AI adoption. We recognise that many institutions struggle to establish consistent policies on assessing algorithmic safety, managing liability, and preserving patient trust. To remedy this, our committees produce recommendations that reflect the latest evidence and are adaptable to diverse regulatory contexts worldwide. We also offer extensive training programmes that cater to professionals seeking practical grounding in AI's potential, limitations, and ethical dimensions. Our two-day introductory course provides an intensive orientation for clinicians who wish to familiarise themselves with AI concepts in a condensed timeframe, and our certificate and diploma tracks, which include case studies and interactive simulations, allow deeper exploration of the subject. We are actively planning the development of a robust MSc programme that will merge academic rigor with clinical applicability, enabling participants to become leaders in the responsible integration of AI across various healthcare domains.

To fulfil these objectives, we have established multiple committees and working groups dedicated to distinct areas of inquiry. Some focus on clinical validation and seek to define standards by which AI tools can be trialled in controlled settings, measured against established diagnostic and therapeutic benchmarks, and audited for ongoing quality once adopted in practice. Others concentrate on ethical and legal frameworks, parsing the complexities of privacy, consent, liability, and bias mitigation. A further group focuses on training and curriculum development, ensuring that each new course or certification meets a consistently high standard so that both novices and experienced professionals gain tangible skill sets. These committees draw their membership from clinicians, computer scientists, ethicists, and legal scholars,

reflecting ISMAI's commitment to robust multidisciplinary dialogue.

By uniting our members' complementary strengths under a single organisational umbrella, we aspire to offer clarity and confidence to the healthcare ecosystem. Hospitals and clinical teams who approach AI with caution will find in ISMAI a partner capable of translating technical jargon into clear guidance, and technology developers will benefit from input grounded in frontline patient care. We see ourselves as not just gatekeepers of safety but also architects of progress who help stakeholders achieve the remarkable benefits that AI can bestow when it is crafted and deployed with fidelity to the mission of healing. Through collaborative research and consensus building, ISMAI strives to shape a medical future in which technology remains an ally rather than an authority and in which human insight remains irreplaceably at the centre.

VI. A Call to Action

In an era of rapid technological change, the future of medicine hinges on an unwavering commitment to patient welfare. The perspectives shared in this white paper reflect a single overarching principle: effective AI must be guided by those who uphold the highest ethical and clinical standards. Voices from diverse fields—physicians, technologists, policymakers, and ethicists—tell us that vigilance alone is not enough. We must be organised, proactive, and united in our resolve to channel AI's potential towards humane, equitable care.

ISMAI was formed precisely to meet that challenge. We centre clinicians in the design, validation, and governance of AI systems because we recognise that these technologies can fulfil their promise only when they amplify rather than erode or replace clinical expertise. ISMAI's role is to create a cohesive framework in which doctors, healthcare institutions, developers, and regulatory bodies coordinate using a shared set of standards. Such coordination is essential if we are to prevent the emergence of a fragmented patchwork of AI solutions with uneven safeguards and uncertain accountability. Our aim is to protect patients while encouraging innovation that genuinely addresses gaps in healthcare delivery.

We invite all professionals and organisations who share this vision to join ISMAI. Whether you are a practicing clinician wishing to bring your frontline perspectives to our committees, a technologist striving to build algorithms that are rigorously validated, or a hospital leader seeking guidance on safe integration, there is a place for you in our society. Membership brings you into contact with peers who are tackling similar challenges, fostering collaboration across disciplinary boundaries and geographical borders. Participation also enables direct engagement with the guidelines and standards we are actively developing to ensure that our collective expertise shapes policy instead of passively reacting to it.

In addition to membership, we offer opportunities to serve on specialised committees that handle the nuanced questions raised by AI's evolving role in medicine. We believe that the tasks of establishing liability, guaranteeing transparency, monitoring biases, and ensuring fair access are so complex that they demand collaborative thinking by practitioners well versed in different domains. We also encourage you to explore our training programmes: a two-day course designed for rapid immersion in AI fundamentals and the more extensive certificate and diploma tracks for people seeking deeper mastery. Soon, we plan to launch an MSc programme that integrates academic rigor with real-world application, allowing participants to learn and

shape the discourse on responsible AI deployment.

This moment requires leadership that is both principled and pragmatic. If we hesitate, governance structures may be devised by those who lack a clear understanding of the ethical duties that guide healthcare. If we remain scattered, we risk repeating past failures. By uniting under ISMAI, we can ensure that AI not only safeguards patient well-being but also upholds a broader social contract that respects autonomy, fairness, and accountability. It is our hope that you will stand with us in this endeavour so we can build a future where medical AI consistently reflects the highest ideals of our profession.