

# INTELLIGENT MEDICINE

A FOUNDING MANIFESTO

---

Nazal Karadan

MD | MSc by Research | MSc

2026

*“Our greatest responsibility is to be good ancestors.”*

— Jonas Salk

## **A NOTE ON THIS MANIFESTO**

---

This is a founding manifesto. So it is not a final word. It makes the case for a new field and describes the governing logic that field requires. This document does not attempt to settle every question the field faces or will face.

If you are a clinician, you will find it useful to understand why the tools you need don't always reach you, and what the systems around those tools must look like before they can.

If you are an engineer or technologist, you will find it useful to grasp that the hardest problem is not always in building the system, but in governing its survival in a real institution.

If you are an investor or a capital allocator, you will find it useful to realize why governance is not always a cost but what makes returns possible, and why the risk profile of clinical systems changes at a fundamental level when they are built within this framework.

If you are an administrator or institutional leader, you will find it useful in finding an honest account of why failed deployments share the same root cause, and what changes when that cause is addressed from the beginning.

Read it in whatever order serves you. The argument is built sequentially, but each section stands on its own.

## THE CASE FOR A NEW FIELD

---

Medicine has been permanently altered. Intelligence — computational, algorithmic, robotic, molecular — is now a ubiquitous part of every layer of clinical practice, from the laboratory to the operating theater, to the clinic, to the ward, to the home. Intelligence in medicine is now almost ambient. This transformation did not come from outside the field, but emerged from its core, reshaping what clinicians can detect, what surgeons can do, what researchers can model, and what patients can expect.

Yet, this transformation has no name, no unified field, nor a governing logic that holds it together across its full scope. We have been calling it by different names – clinical AI, surgical robotics, digital therapeutics, computational imaging, continuous sensing, precision pharmacology. Each is treated as its own vertical, its own discipline, its own problem to solve in isolation. The result is a field of extraordinary parts with no coherent whole.

Intelligent Medicine is the name for it. It is the name for what medicine is already becoming. We now have sufficient knowledge and foresight to consider it a field in its own right. And it is the governing field that ensures that becoming is purposeful, integrated, durable and equitable, rather than merely accidental, fragmented, episodic, and exclusive.

The case for naming it is not abstract. Intelligent technology has already entered medicine without a unified field to receive it. And the consequences of that absence are visible everywhere, but nowhere more clearly than in the distance between discovery and delivery, one of medicine's oldest structural failures. Across clinical environments and research institutions, this pattern appears with remarkable consistency.

Penicillin, for instance, was discovered in 1928. But it did not reach patients at scale until the infrastructure of war forced it into existence well over a decade later. In 1847, Semmelweis demonstrated that handwashing prevented deaths on maternity wards. The institutions of the time destroyed him for it. This pattern predates computation, digitization, and artificial intelligence by centuries.

With intelligent technology, what has changed is not the existence of that gap, but the scale and speed, accelerating the distance between innovation and delivery.

Somewhere right now, a diagnostic tool that could help a clinician detect early-stage disease is sitting in a research paper. It will not reach a patient this year, as no system has been built around it, nor reimbursement modeled, nor regulatory pathway addressed, nor the

clinical workflow redesigned to accommodate it. These are not isolated events and they are often brushed off as the default outcome.

But this gap has a human cost that can be measured in outcomes, an economic cost that can be measured in capital deployed against systems that were never designed to endure, as well as an institutional cost, measured in the erosion of trust each failed deployment leaves behind.

*Intelligent Medicine is the field that governs the convergence of biological, computational, and institutional systems in clinical environments — and it is built to close the gap.*

## **I. THE STRUCTURAL BREAK**

---

Healthcare is no longer constrained by intelligence. It is constrained by its structure.

Algorithms now enhance clinical decision-making in domains that previously would have required several years of specialist training to navigate. Robotic systems extend surgical precision beyond the natural limits of the human hand. Sensors stream continuous, real-time physiological data across populations which would have been impossible with manual monitoring. Molecular diagnostics help compress biological complexity into actionable clinical signals.

Biological AI, where computational models that predict molecular behaviour, protein structure, and genomic function now work at a scale no human team can match. It is compressing the distance between biological complexity and clinical application.

Despite all of this, hospitals remain fragmented. Breakthroughs and technologies do not consolidate or endure. The same structural failures repeat for decades across institutions and disciplines.

Today medicine has acquired an entirely new class of intelligent tools without developing the governing field required to integrate them. And while these tools are unprecedented, the governing logic has not arrived to match them.

*Intelligent Medicine names that governing logic, not as a metaphor or as a framework applied within existing fields, but as a field in its own right.*

## II. A FIELD, NOT A FRAMEWORK

---

This distinction matters greatly. A framework is a tool someone uses, whereas a field is a discipline someone belongs to. A discipline has its own object of study, its own body of knowledge, its own criteria for success, and its own accountability for outcomes. Intelligent Medicine is a field. Not a framework borrowed from medicine or engineering or economics, but a discipline that sits at their convergence and governs what none of them can govern alone.

Clinical AI, surgical robotics, biotech, digital therapeutics, computational imaging, and continuous sensing are currently treated as separate verticals or specialties. They have their own conferences, investment theses, talent pipelines, and governance problems. Yet, they are all expressing the same underlying transition: medicine becoming an intelligent system. As a result, they share the same structural failures and require the same governing logic to survive.

What Intelligent Medicine studies and governs is the convergence of biological, computational, and institutional systems in clinical environments. That convergence does not belong to medicine alone, or to engineering, or to economics. It belongs to Intelligent Medicine.

Intelligent Medicine is not artificial intelligence applied to healthcare, nor is it digital health, eHealth, or any of the predecessor categories that came before. But it incorporates what was useful in each and discards the noise in each. It is also not a collection of smart devices or a platform or a product line. Each of those categories addresses one thing. Intelligent Medicine addresses the whole.

*Intelligent Medicine is what turns fragmentation into a field.*

## III. THE TRANSLATION PROBLEM

---

The failure pattern is consistent enough to be predictable. Anyone who has watched a promising technology's attempt to make it to real-world hospital systems will recognize it almost immediately.

We regularly see how a system can perform well under controlled conditions but still degrade under real-world complexity. Or, how a regulatory pathway is treated as an endpoint rather than a design constraint, killing its survival in the system. Or, failure via

reimbursement models that are never built, or clinicians stopping using a tool within weeks because the tool adds friction via time pressure or liabilities, rather than removing it.

All of these are considered norms in the industry, and most of them are not unpredictable either. In fact, they are the direct consequences of building systems without thinking about how to embed the conditions for their survival from the very beginning.

The fragility we observe is by design, for a system that is validated only in a controlled environment is bound to be fragile. Without economic alignment, no technology can scale, and a technology will most likely be deemed useless if it increases cognitive load.

Intelligent Medicine considers the full picture, finds solutions, and embeds antidotes from inception. This includes validation under variability, regulatory architecture as a design input, building economic alignment into the development model, and strong workflow integration.

*If a system cannot survive real-world constraints, it is not ready to become medicine.*

#### **IV. WHAT GOVERNANCE ACTUALLY LOOKS LIKE**

---

We cannot merely rest on abstract principles. The value of Intelligent Medicine as a field lies in what it changes in practice.

Consider a clinical AI system that is designed to enhance sepsis detection in ICU patients, where the algorithm is validated, the metrics are strong, the publication is peer-reviewed. Under conventional development, that would be considered rigorous, but it is also where the rigor stops. But under Intelligent Medicine, it is where the real work begins.

Under Intelligent Medicine, the system would have to be validated not just on curated datasets but under the heterogeneity of real ICU populations of different institutions, equipment, and documentation practices. The regulatory pathways have to be mapped before products are built, as it is the architecture of compliance that shapes the architecture of the system. Reimbursement models have to demonstrate that the system is cost-effective or improves the outcome in such a way that the payers recognize the change. Finally, the clinical workflows have to be redesigned around the tool, rather than a tool that is inserted into an existing workflow, creating friction.

And while none of this is innovation, it all determines whether the innovation survives.

Very often the adoption of clinical technology is treated as a procurement decision. Its governance, including accountability, validation, and workflow redesign, is addressed after the contract is signed, if at all. This is a consequence of not having a field responsible for asking the governance questions before the purchase order is raised.

As clinical AI moves from monolithic models toward modular agent-based architectures, this governance challenge deepens. Validating a system is no longer a single event, rather it becomes a continuous discipline applied across interacting agents, orchestration layers, and decision pathways that evolve in deployment. And traceability becomes the mechanism through which institutional trust is built over time.

Governance also requires defining accountability. When a modular system makes a decision that affects a patient, the question of responsibility is structural rather than merely technical. And so we must ask: who is accountable? The developer, the institution, the clinician who relied on it, or the governance framework that certified it?

*Intelligent Medicine must define accountability before deployment.*

The interface between human clinical judgment and autonomous system output is not a philosophical boundary. It is a legal and institutional one that determines whether any system can be sustainably deployed at scale.

## **V. USEFULNESS, EFFICIENCY, AND EQUITY**

---

Three principles anchor the field.

The first is usefulness. An intelligent system must alter clinical reality in measurable terms, whether that is enhancing diagnostic precision, reducing preventable risk, improving safety, or optimizing therapeutic outcomes. It is not sufficient to demonstrate this capability. Utility must withstand real-world scrutiny, and not just controlled evaluation.

The second is efficiency. Medicine operates within hard constraints of time, staffing, attention, and capital. An intelligent system must strengthen this environment, not add to its burden. It must reduce cognitive load, compress workflow wherever possible, and protect clinical throughput. Precision without usability is ornamental; intelligence without integration is wasteful.

The third is equity. Intelligent technology, properly governed, has the potential to actively reduce health inequality. A system that improves diagnostic accuracy in a well-resourced teaching hospital but cannot be deployed in a district clinic has not advanced medicine. It has advanced the advantage of those who already have it. Intelligent Medicine must govern systems deployable across resource settings. If a governance framework only works for the best-funded hospitals, it has not solved the translation problem.

Equity is not an abstract principle. In practice it means designing for the oncologist managing two thousand patients, not just the one managing two hundred. It means building systems that function in a seven-minute consultation, and not just in a well-resourced multidisciplinary team meeting. It means ensuring that a patient receiving a cancer diagnosis can understand their own report without needing a medical degree. These are the median reality of medicine globally. Intelligent Medicine is not a premium product dressed as a solution, but one built for the reality of global medicine.

There is also an economic argument. Intelligent clinical systems, properly governed and deployed at scale, have the potential to reduce the cost of care delivery dramatically. It does so by compressing the gap between what quality medicine costs today and what it needs to cost to reach everyone who needs it. That is not idealism, but the same logic that makes any mature technology cheaper as governance matures around it. Medicine is not exempt from that trajectory.

*Usefulness establishes legitimacy. Efficiency determines survival. Equity determines whether the field is building for people or privilege.*

## **VI. REGULATION AND CAPITAL AS DESIGN FORCES**

---

Healthcare is an institutional system that is governed by law, economics, and risk. It is the very operating environment, the environment technology must be designed for.

Regulatory frameworks are routinely treated as obstacles, or stages that must be passed after the real work is done. This is a category error. Intelligent systems are being developed and deployed faster than regulatory frameworks can track in isolation. The answer is not to slow the systems or ignore the frameworks, but it is to continuously build the relationship between the regulators and Intelligent Medicine practitioners. They must work together from the beginning of development, not at the point of submission. When that happens,

regulation becomes a net positive, structural filter that strengthens systems rather than a bureaucratic barrier that delays them.

*Regulation, integrated architecturally from the beginning, becomes a competitive advantage.*

Capital follows the same logic. When regulation and deployment are disciplined and work together, the risk profile of intelligent clinical systems changes fundamentally. Capital does not fund perpetual experimentation. Instead it funds systems that can scale without destabilizing the institutions that carry them. When governance is rigorous and deployment is more predictable, capital flows because the uncertainty has been structurally reduced.

A clinical system that survives these constraints is not merely a technological achievement. It is a durable institutional asset, and one that can be defended to regulators, justified to boards, and sustained through the ordinary turbulence of institutional life.

*Medicine that cannot sustain itself economically cannot persist clinically.*

## **VII. THE GOVERNANCE LAYER MEDICINE IS MISSING**

---

The fragmentation of medicine is not accidental. But it reflects a historical fact about the field, on how it has traditionally allocated responsibilities.

It is standard that clinicians focus on care, engineers on tools, regulators on safety, administrators on budget, and investors on capital returns. On the surface it is a fairly neat distribution of tasks, where each domain has developed deep expertise within its own boundaries, and none of them are responsible for what happens across them.

But as technology accelerates, this fragmentation becomes a central problem. The complexity of modern clinical systems exceeds what any single domain can govern. All decisions that determine whether the system succeeds or fails are not contained within clinical medicine, engineering, regulation, or finance. Instead they live and breathe at the intersections. The technology exists at these intersections and they have no owner.

*Intelligent Medicine introduces the missing governance layer.*

This is leadership capable of integrating biological complexity, computational architecture, regulatory terrain, and institutional economics into a coherent system. It is not a committee or a consulting entity. It is a field, with its own logic, where practitioners are

trained across domains, not just within them, so that they understand the complexity of the system, and are accountable for outcomes across the system.

Institutions that develop this capacity will be able to do something their competitors cannot: take the tools that exist right now and actually build something durable with them. Institutions that do not will continue to accumulate impressive technologies and wonder why nothing transforms.

Building this capacity requires people who can operate across the full stack. Three gaps stand out.

The first is the clinician with genuine AI literacy. This is someone with an operational understanding of how intelligent systems are built, validated, and governed. A clinician who can evaluate an algorithm's clinical claims with rigor. This person does not currently exist in sufficient numbers because medical training has never required it. The tools arrived before the curriculum did, but that gap is not abstract. We see it show up every time a hospital committee evaluates a clinical AI system and nobody in the room can interrogate the validation methodology.

The second is the regulatory scientist trained for dynamic systems. Current regulatory professionals are built around static device frameworks. They deal with the approval and deployment of a device, which stays the same for a long period. Clinical AI systems learn, drift, and update. They also behave differently across populations. They fail in ways that static frameworks were never designed to handle. The field needs practitioners who understand continuous validation, post-market surveillance for adaptive systems, and ongoing governance obligation.

The third is the clinical operations leader capable of governing the full deployment lifecycle. This is not a data scientist or a pure clinician or an administrator, but someone who understands all three well enough to make a system survive contact with a real institution. This person understands the workflows, the politics, the resource constraints, and the tolerance for change in the institution. This is the leadership role that determines whether a validated, approved, economically aligned system actually gets used. It is the most critical role in the ecosystem and the most consistently absent from organizational charts.

Intelligent Medicine is the field responsible for producing all three. Its emergence as a discipline is inseparable from its emergence as a profession.

That emergence requires institutional commitment. This could mean designing a curriculum that does not exist yet, including training programmes that cut across clinical medicine, engineering, regulatory science, and institutional economics simultaneously. It means creating roles within hospitals, health systems, and medtech companies that have not been formally defined. It means building the academic and professional infrastructure that transforms a founding argument into a working field. None of this happens automatically. Institutions will need to decide to build it, and it will require enough people who understand the logic of converging fields.

## **VIII. DATA, SOVEREIGNTY, AND THE PATIENT**

---

Intelligent systems run on data. The governance of those systems is inseparable from the governance of the data that drives them.

Clinical data is however not a neutral resource. It is generated by patients, carries their biological and personal identity, and holds commercial value that grows as systems scale. Clinical data raises vital governance questions: who owns it, who controls access to it, and how the value it generates is distributed. These questions determine whether it serves patients or extracts from them.

Intelligent Medicine must govern the data layer as rigorously as it governs the clinical layer. This means embedding data sovereignty into system architecture from the inception — not merely as a compliance requirement but as a structural commitment. Patients must understand what their data is used for, how it influences the systems that affect their care, and what protections exist when those systems fail.

*The patient is not merely the endpoint of intelligent medicine. They are a participant in its architecture.*

## **IX. WORKING WITHIN INSTITUTIONS, NOT AGAINST THEM**

---

Intelligent Medicine does not propose to replace existing medical institutions or systems. It proposes to change what they are capable of.

Hospitals, health systems, and regulatory bodies are not obstacles. But they are the environment in which clinical capability and intelligent technology must survive. The

question is not how to circumvent them, but how to build systems that can operate within them without being diminished by them.

What Intelligent Medicine challenges is not the institution but the assumption that institutional structure and technological advancement have an inherent tension between them. In reality, they do not. The tension exists because the governance layer that would allow them to operate together has been missing. People and systems adapt faster than institutional pessimism suggests. What they need is not more time, but the right framework. And that is what this field builds.

If you are a clinician who sees what the tools could do and is frustrated by what the systems prevent, this field is built on that frustration. If you are an engineer who builds what works in the lab and watches it fail in the institution, this field addresses that failure. If you are an investor who has funded the technology and cannot understand why it does not scale, this field is the answer to that question. If you are an administrator who inherits the consequences of every failed deployment, this field is what improves the probability of the next one.

The problem is shared. So is the cost. This field is the framework for addressing both.

## **X. THE LIMITS OF THE FIELD**

---

No field is immune to the failure it was designed to prevent. Intelligent Medicine is not exempt from this.

The risk is familiar: a field created to reduce institutional friction becomes institutional friction itself, or a governance layer introduced to accelerate integration slows it instead. The framework becomes a checklist, and the checklist becomes a bottleneck. The discipline meant to serve clinical outcomes begins to serve its own perpetuation. It must be challenged and reformed when it does, because the aim of Intelligent Medicine is structural maturation, not institutional self-preservation.

The check against this failure is the same one the field applies to everything it governs — usefulness and efficiency. If the governance layer cannot demonstrate that it improves outcomes and reduces waste, it has failed on its own terms.

The concrete version of this failure looks like a validation committee that adds six months to every deployment regardless of system complexity. Or a compliance framework

that was designed for hardware devices and is now applied unchanged to adaptive software, producing requirements that are technically satisfied but clinically meaningless. Or a governance body whose primary output is documentation rather than decisions. These are the predictable endpoints of any governance system that stops asking whether it is working and starts assuming that its existence is sufficient.

The field must be held to the standard it sets for the tools it governs. Intelligent Medicine that does not make medicine more intelligent is not Intelligent Medicine. It becomes bureaucracy with better branding.

## **XI. WHAT MEDICINE LOOKS LIKE WHEN IT WORKS**

---

A healthcare system that has internalized Intelligent Medicine does not look radically different from the outside. The hospitals and the physicians are still there, but something fundamental has changed in how the system learns, adapts, and sustains itself. The flow of the system is different. It is faster where speed saves lives, more deliberate when deliberation is required.

Technologies reach patients because the pathway was specifically designed for it. From its inception, a diagnostic tool validated in research carries with it the regulatory architecture, the economic model, and the workflow integration that allow it to survive contact with clinical reality. And deployment becomes the final stage of a process that began at the drawing board.

Clinicians will operate with tools that reduce their burden rather than add to it, because they were designed by people who understood medicine, and understood that a tool which increases cognitive load will always lose to the demands of a ward at capacity.

Capital flows toward systems rather than devices. Systems where data serves patients who are no longer passive participants, because their data is governed and not merely extracted.

*The gap between what medicine knows and what medicine delivers begins to close systematically rather than accidentally.*

## **XII. THE HORIZON**

---

The trajectory is already visible. Intelligent agents will coordinate triage and continuous monitoring across patient populations too large for manual oversight. Sensor technologies and wearables will extend clinical care beyond hospital walls into the rhythms of daily life. Robotic systems will help execute interventions informed by predictive analytics updated in real time. Therapeutics will adapt dynamically to molecular and physiological feedback, moving medicine from population-level approximation towards real biological precision. Biological AI is accelerating that precision further by moving drug design and molecular diagnostics from broad population assumptions toward individual biological reality.

At the population scale, intelligent infrastructure will change how medicine responds to systemic threats. Pandemic preparedness, outbreak surveillance, and global health response have historically been reactive, often mobilized after the crisis has already spread. The same infrastructure that monitors an individual patient can, at scale, detect population-level anomalies before they become emergencies. Intelligent Medicine in public health is about seeing what is coming before it arrives.

As technical sophistication increases, the rate of invention will not slow. Without structural integration, fragility will increase with it.

The advantage in the next era of healthcare will not belong to those who build the most advanced tools. It will belong to those who architect the most resilient systems, those who understand that the core limiting variable is not invention but integration.

*Integration will define competitive power. Durability will define leadership. The future belongs to institutions and systems that can govern convergence.*

This manifesto exists because the gap between what medicine can build and what it can deliver is in front of us and one we can no longer ignore.

Intelligent Medicine is the field that ensures intelligence becomes infrastructure, that the tools being built right now evolve into validated, regulated, adopted, and economically sustainable clinical capability, and matures towards becoming a part of the health systems we work in.

*Intelligent Medicine is the new ground we have been waiting to build on.*

## **ABOUT THE AUTHOR**

---

Nazal Karadan is a physician-scientist working at the intersection of clinical medicine, neuroscience, bionic engineering, and systems thinking. He is an MD with a master's by research in neuroscience from the University of Edinburgh, and a master's in bionic engineering from Germany.

Intelligent Medicine was not developed in response to a trend. It was developed in response to a structural failure observed across institutions, disciplines, and markets — one that existing fields were not positioned to address. The manifesto is a founding document, and not the final word.

*If you would like to translate this manifesto into another language, you are welcome to do so with attribution.*

*First Edition | 2026*