MAR Oncology and Hematology (2026) 6:3

Research Article

From Atoms to Agents: A Decade of Transformation in Artificial Intelligence (AI) in Medicine

Dr. Enrique Díaz Cantón, MD, MSc, MSc *

***Correspondence to**: Dr. Enrique Díaz Cantón, MD, MSc, MSc, Professor of Oncology and Artificial Intelligence in Medicine, CEMIC University Institute, Buenos Aires, Argentina.

Director of the Postgraduate Program in Artificial Intelligence and Medicine, National Academy of Medicine, Buenos Aires, Argentina

Received: 07 March 2026

Published: 01 April 2026

DOI: https://doi.org/10.5281/zenodo.19353904

Abstract

The past decade has witnessed an unprecedented acceleration in the capabilities of artificial intelligence (AI), catalyzing a paradigm shift in medicine. This special article synthesizes a series of recent investigations across three key evolutionary stages of this transformation. First, it analyzes the transition from static structural biology to dynamic molecular simulation, driven by an ecosystem of models such as AlphaFold, which predicts protein structure; AlphaFold 3 and AlphaMissense, which interpret molecular interactions and genomic variants; and the Isomorphic Drug Design Engine (IsoDDE). This ecosystem, further complemented by the emergence of World Models, is redefining rational drug discovery and ushering in the era of computational oncology. Second, the article explores the "agentic shift," in which AI evolves from a passive analytical tool into a proactive, semi-autonomous collaborator. The rise of multi-agent AI systems, orchestrated through frameworks such as the Model Context Protocol (MCP), promises to reshape clinical workflows, from multidisciplinary tumor boards to adaptive radiotherapy planning. Finally, it addresses the profound ethical, philosophical, and regulatory implications of this progress. In particular, it examines the future of medical work, the philosophical debate surrounding consciousness and subjective experience in AI, and the urgent need to establish clear legal and ethical frameworks, as illustrated by the evolving regulatory landscape in Argentina. This synthesis offers a comprehensive perspective on the trajectory of AI in medicine, from the atomic precision of molecular modeling to the collaborative complexity of agentic systems and the humanistic challenges that lie ahead.

Keywords: *Artificial Intelligence; Computational Oncology; Multi-Agent Systems; Medical Ethics; Drug Design; World Models.*

Introduction

Artificial intelligence (AI) has ceased to be a futuristic promise and has become a transformative force in the present of medicine. Over the course of the last decade, we have witnessed a rapid evolution that has taken AI from the interpretation of images and tabular data to the ability to predict the three-dimensional structure of biological macromolecules, simulate their dynamic behavior, and, more recently, function as a proactive collaborator in complex clinical decision-making. The magnitude of this transformation is illustrated by the exponential growth in AI-based medical devices authorized by the FDA, which surpassed 1000 cumulative approvals in 2024 [1], and by the inclusion of the first predictive AI tool in oncology clinical practice guidelines [2].

This special article summarizes a series of recent works that trace this evolutionary trajectory and argues that we are currently undergoing a transition on three major fronts: (1) the leap from static structural modeling to dynamic simulation for rational drug design, (2) the emergence of multi-agent AI systems that are redefining human-machine collaboration in the clinical environment, and (3) the rise of unavoidable ethical, philosophical, and regulatory questions regarding the future of medical work and the nature of intelligence itself.

From Structure to Dynamics: The AI Ecosystem for Drug Discovery

The inflection point in modern computational biology came with AlphaFold 2, a deep-learning system capable of predicting the three-dimensional structure of proteins from their amino acid sequences with near-experimental accuracy [3]. This achievement, recognized with the 2024 Nobel Prize in Chemistry, solved a more than fifty-year-old scientific challenge and generated a public database containing predicted structures for more than 200 million proteins [4]. Yet this structural representation—analogous to a molecular "snapshot" frozen in time—omits the fundamental dimension of protein dynamics: conformational fluctuations, transition states, and allosteric movements that are critical for biological function and drug interaction.

Precision oncology, in particular, has long faced the barrier of so-called "undruggable" therapeutic targets, such as oncogenic mutations in KRAS, p53, or MYC, often because of the absence of well-defined binding cavities in their static state. The solution to this impasse is now emerging from an integrated AI ecosystem that overcomes the limitations of static modeling.

AlphaFold 3 extended the capabilities of its predecessor by modeling not only isolated proteins but also molecular complexes that include nucleic acids, ligands, and ions, thereby providing a more realistic view of the biomolecular environment [5]. In parallel, AlphaMissense made it possible to classify the pathogenicity of nearly all possible missense variants across the human proteome, providing a genome-scale tool for the interpretation of variants of uncertain clinical significance [6]. More recently, the drug design engine developed by Isomorphic Labs (IsoDDE) represented a qualitative advance by going beyond structural prediction alone. IsoDDE doubles the accuracy of AlphaFold 3 in predicting out-of-distribution protein–ligand structures, outperforms physics-based benchmark methods (such as FEP+) in predicting binding affinity, and identifies cryptic binding pockets—including novel allosteric sites in cereblon—using only the amino acid sequence as input [7]. Crucially, IsoDDE models "induced fit" phenomena, in which the protein changes its conformation to accommodate the ligand, suggesting the emergence of implicit world models within molecular prediction.

The next conceptual frontier is the explicit adoption of "World Models"—a paradigm originally proposed by Ha and Schmidhuber [8] and further expanded in architectures such as Genie by Google DeepMind and JEPA by Meta AI [9,10]—to learn the dynamic rules that govern the evolution of a biological system. Applied to molecular dynamics, these models promise not merely to predict a final structure, but to simulate the conformational trajectory of a protein in response to a stimulus such as drug binding or a point mutation. This opens the door to the rational design of allosteric modulators and to the creation of therapies directed against intrinsically disordered proteins, marking the true dawn of computational oncology [11,12].

The Emergence of Multi-Agent Systems in Clinical Practice

If the first phase of the AI revolution in medicine focused on data analysis and pattern classification, the current phase is defined by the capacity for autonomous action. We are witnessing an "agentic shift," in which AI is moving from being a passive tool that responds to queries to becoming an active, semi-autonomous teammate capable of planning, executing sequences of tasks, and using external tools [13,14]. This transition is being enabled by the convergence of large language models (LLMs) with agentic frameworks and interoperability protocols.

A key catalyst of this convergence is the Model Context Protocol (MCP), an open standard introduced by Anthropic in November 2024 that provides a universal interface through which AI models can connect to external data sources and software tools [15]. Rapidly adopted by OpenAI, Google DeepMind, and Microsoft, MCP addresses the combinatorial integration problem of linking multiple models to multiple tools, thereby providing the infrastructure required for multi-agent systems.

Rather than relying on a single monolithic model, multi-agent AI systems consist of multiple specialized agents that collaborate to solve complex problems. In oncology, this has transformative implications that have already begun to materialize. Ferber et al. showed that an autonomous GPT-4–based clinical agent, integrated with computational pathology tools, radiologic segmentation systems, and oncology knowledge bases, improved clinical decision accuracy from 30.3% (standalone model) to 87.2% (integrated agentic system) across 20 multimodal oncology cases [16]. Similarly, recent reviews in Nature Reviews Cancer and Federal Practitioner have thoroughly documented how AI agents are emerging in oncology research and clinical practice, from autonomous drug design to the proposal of therapeutic strategies based on the integration of multi-omic data [13,17].

A multi-agent system could, for example, orchestrate the workflow of a multidisciplinary tumor board, in

which one agent retrieves and synthesizes the patient's medical history, another analyzes radiologic images using segmentation models, a third interprets genomic and pharmacogenomic data, and an "orchestrator" agent integrates the information to present a structured summary and evidence-based recommendations to the medical team. Microsoft has already developed a prototype agent orchestrator for oncology management (Healthcare Agent Orchestrator) that integrates multimodal pathology, radiology, and clinical-data models into tumor-board workflows [18].

This agentic capability is already demonstrating clinical value in radiation oncology. Auto-segmentation of organs and tumors through deep learning—one of the most mature AI applications in this field—dramatically reduces planning times from hours to minutes, with a consistency that matches or exceeds interobserver variability among specialists [19,20]. Predictive AI, such as the ArteraAI platform, is already being used to personalize treatment in prostate cancer and is the first AI-based test incorporated into the NCCN guidelines, with level IB evidence [2,21]. Integrating these components into an agentic system would enable real-time adaptive radiotherapy by adjusting treatment plans on the basis of the patient's daily anatomic changes detected on verification imaging.

For regions such as Latin America, which face a substantial shortage of radiotherapy equipment and specialized personnel—with fewer than half the linear accelerators recommended by the World Health Organization—agentic AI represents a singular opportunity to democratize access to high-quality cancer care by reducing dependence on specialist availability at each step of the workflow [22].

The Future of Medical Work: Ethical, Philosophical, and Regulatory Considerations

The exponential advance of AI compels us to confront fundamental questions about our profession, our role as physicians, and the very nature of intelligence itself. As AI models expand their competence from the domain of language (logos) into visual perception, planning, and interaction with the physical world—what some have called "spatial intelligence" and "embodied intelligence"—the debate on consciousness and subjective experience in machines becomes unavoidable [23,24].

From a philosophical perspective, even if an AI were to achieve superhuman competence across all cognitive and procedural tasks in medicine, it might still lack subjective experience. The concept of the "philosophical zombie," articulated by Chalmers in his formulation of the "hard problem of consciousness," raises the theoretical possibility of a system that is functionally indistinguishable from a conscious being yet entirely devoid of qualia, the subjective qualities of experience [25]. This distinction is critical in medicine. An LLM

can process the term “cancer pain” and access the full relevant literature on its pathophysiology, pharmacology, and management, but it does not “understand”—in the phenomenological sense—the suffering that such pain entails. In Nagel’s terms, there is no “what it is like” of pain within its processing [26].

Accordingly, although AI will increasingly assume responsibility for the analysis of massive datasets and the execution of standardized procedures, the physician’s irreplaceable role will reside in the domain of shared experience: empathy, situated ethical judgment, the communication of uncertainty, and the capacity to make sense of illness together with the patient. The physician-patient relationship—as an encounter between two subjectivities—remains an irreducibly human territory, at least within current computational architectures.
This new paradigm also requires a robust and updated regulatory framework. Globally, the FDA has authorized more than 1000 AI-based medical devices, predominantly in radiology, yet without specific legislation for agentic systems that make sequential decisions [1]. In Europe, the European Union’s Artificial Intelligence Act (AI Act), which entered into force in 2024, classifies AI systems in health care as “high risk” and establishes requirements for transparency, human oversight, and conformity assessment [27].

In Argentina, ANMAT has begun to address this need by classifying Software as a Medical Device (SaMD), including AI-based systems, and subjecting them to an approval process based on risk classification according to the guidelines of the International Medical Device Regulators Forum (IMDRF) [28]. In 2019, the first AI software for clinical use in Latin America (Entelai Pic, for neuroimaging) was approved, and more recently AI tools for breast cancer detection have also received ANMAT authorization [29]. However, substantial uncertainty remains regarding legal liability in the event of error by an autonomous AI system, particularly in the context of agentic systems that execute chains of decisions. It is imperative that Argentina and the region develop specific legislation establishing a clear liability framework, mechanisms for algorithmic auditing, and training programs for health professionals so that these tools can be used safely, ethically, and effectively.

Key Points

Current knowledge: AI in medicine has progressed from the analysis of tabular data and the classification of images to the prediction of static molecular structures and assistance in specific tasks such as medical-image segmentation.

Contribution of this article: This paper synthesizes the recent evolution of AI in medicine across three interconnected fronts: (1) the transition from static models to dynamic models and integrated drug design engines (the AlphaFold–IsoDDE–World Models ecosystem), which inaugurate computational oncology; (2)

the emergence of multi-agent systems that, orchestrated through protocols such as MCP, act as semi-autonomous clinical collaborators with growing evidence of efficacy; and (3) the ethical, philosophical, and regulatory challenges—from the problem of consciousness to the need for specific legal frameworks—that define the future role of the physician in an era of advanced artificial intelligence.

Conflict of Interest

The author declares that there are no conflicts of interest.

Declaration on the Use of Artificial Intelligence

Generative artificial intelligence tools (Claude, Anthropic) were used as assistants during the literature search, the structural organization of the content, and the preliminary drafting of this manuscript. All content was fully reviewed, verified, edited, and validated by the author, who assumes full responsibility for the accuracy, originality, and intellectual integrity of the final work. References were individually verified by the author. The use of AI is disclosed in accordance with current ICMJE recommendations regarding transparency in the use of these technologies in scientific production.

References

1. Muehlematter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–2024): a comparative analysis. Lancet Digit Health. 2024;6(6):e449-e458.

2. Esteva A, Feng J, van der Wal D, et al. Prostate cancer therapy personalization via multi-modal deep learning on randomized phase III clinical trials. NPJ Digit Med. 2022;5(1):71.

3. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596(7873):583-589.

4. Varadi M, Anyango S, Deshpande M, et al. AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. Nucleic Acids Res. 2022;50(D1):D439-D444.

5. Abramson J, Adler J, Dunger J, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024;630(8016):493-500.

6. Cheng J, Novati G, Pan J, et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. Science. 2023;381(6664):eadg7492.

7. Isomorphic Labs. The Isomorphic Labs Drug Design Engine unlocks a new frontier beyond AlphaFold. Technical report. February 2026. Available from: https://www.isomorphiclabs.com/articles/the-isomorphic-

labs-drug-design-engine-unlocks-a-new-frontier

8. Ha D, Schmidhuber J. Recurrent World Models Facilitate Policy Evolution. In: Advances in Neural Information Processing Systems 31 (NeurIPS 2018). 2018.

9. Bruce J, Dennis M, Edwards A, et al. Genie: Generative Interactive Environments. In: Proceedings of the 41st International Conference on Machine Learning (ICML 2024). 2024.

10. LeCun Y. A path towards autonomous machine intelligence. OpenReview [preprint]. 2022. Available from: https://openreview.net/pdf?id=BZ5a1r-kVsf

11.Díaz Cantón E. From Protein Folding to Rational Drug Design: The DeepMind–Isomorphic Labs Ecosystem and the Dawn of Computational Oncology. Int J Tumor Res. Epub ahead of print February 2026.

12. Díaz Cantón E. The Fusion of Artificial Intelligence and Omics: A Perspective Toward 2035. Revista de la Sociedad Argentina de Oncología Clínica. In press, 2026.

13. Truhn D, Akhoundova D, Liechti R, et al. Artificial intelligence agents in cancer research and oncology. Nat Rev Cancer. 2026;26(3):145-162.

14. Gao S, Fang A, Huang Y, et al. Empowering biomedical discovery with AI agents. Cell. 2024;187(22):6125-6151.

15.Anthropic. Introducing the Model Context Protocol. 2024. Available from: https://www.anthropic.com/news/model-context-protocol

16. Ferber D, Kühn D, Bever T, et al. Development and validation of an autonomous artificial intelligence agent for clinical decision-making in oncology. Nat Cancer. 2025. doi:10.1038/s43018-025-00991-6.

17. Dinh-Le C, Lehman R, Pebdani RN, et al. Multiagent AI Systems in Health Care: Envisioning Next-Generation Intelligence. Fed Pract. 2025;42(5):188-194.

18. Microsoft Research. Developing next-generation cancer care management with multi-agent orchestration. Microsoft Industry Blog. May 2025. Available from: https://www.microsoft.com/en-us/industry/blog/healthcare/2025/05/19/developing-next-generation-cancer-care-management-with-multi-agent-orchestration/

19. Rong Y, Chen Q, Fu Y, et al. NRG Oncology assessment of artificial intelligence deep learning-based auto-segmentation for radiation therapy: current developments, clinical considerations, and future directions. Int J Radiat Oncol Biol Phys. 2024;119(1):261-280.

20. Kawula M, Purice D, Picciotto G, et al. Deep learning for autosegmentation for radiotherapy treatment planning: state-of-the-art and novel perspectives. Strahlenther Onkol. 2024;200(9):753-772.

21. Spratt DE, Tang S, Sun Y, et al. Artificial intelligence predictive model for hormone therapy use in prostate cancer. NEJM Evid. 2023;2(8):EVIDoa2300023.

22. Díaz Cantón E. Artificial Intelligence in Radiation Oncology: From Promise to Clinical Practice—A

Transformative Opportunity for Latin America. Revista de la Sociedad Argentina de Terapia Radiante Oncológica. In press, 2026.

23. Li F-F. Spatial Intelligence. TED Talk. 2024. Available from: https://www.ted.com/talks/fei_fei_li_with_spatial_intelligence_artificial_intelligence_will_understand_the_real_world

24. Díaz Cantón E, Casarini I. In Ipso Vivimus et Movemur: Spatial Intelligence, World Models, and the Threshold of Embodied Artificial General Intelligence. Iberoamerican Journal of Medicine. Epub ahead of print 2026.

25. Chalmers DJ. Facing up to the problem of consciousness. J Conscious Stud. 1995;2(3):200-219.

26. Nagel T. What is it like to be a bat? Philos Rev. 1974;83(4):435-450.

27. European Parliament and Council. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. 2024.

28. ANMAT. Software as a Medical Device (SaMD). Public consultation document. 2024. Available from: https://opinionpublica.anmat.gob.ar/proyectos/5293.pdf

29. Narvaez D, Díaz Cantón E. Regulation of AI in Medical Imaging and Diagnostics in Argentina: A Regulatory and Clinical Overview. Technical report. Asociación Argentina de Oncología Clínica. 2026.