Research Article

Journal of Artificial Intelligence in Healthcare & Medicine

Convergence of Large Language Models and World Models in drug Discovery: from the AlphaFold–Isomorphic Ecosystem to Molecular World Models in HER2-Positive Breast Cancer

Enrique Díaz Cantón, MD, MSc, MSc*

National Academy of Medicine of Buenos Aires, Department of Clinical Oncology, Postgraduate Programme in Artificial Intelligence and Medicine, Instituto Universitario CEMIC, Buenos Aires, Argentina.

***Correspondence:**
Dr Enrique Díaz Cantón, National Academy of Medicine of Buenos Aires, Department of Clinical Oncology, Postgraduate Programme in Artificial Intelligence and Medicine, Instituto Universitario CEMIC, Buenos Aires, Argentina.

Received: 09 May 2026; **Accepted:** 10 Jun 2026; **Published:** 21 Jun 2026

Citation: Enrique Díaz Cantón. Convergence of Large Language Models and World Models in drug Discovery: from the AlphaFold–Isomorphic Ecosystem to Molecular World Models in HER2-Positive Breast Cancer. J Artif Intell Healthcare Med. 2026; 1(1); 1-5.

ABSTRACT

Drug discovery is undergoing a historic transition driven by artificial intelligence. The DeepMind ecosystem — AlphaFold 2 and 3, AlphaMissense, AlphaGenome— and its therapeutic arm, Isomorphic Labs, have moved the field from static structural prediction to the rational design of ligands; the Isomorphic Drug Design Engine (IsoDDE), unveiled in February 2026, more than doubles the accuracy of AlphaFold 3 on out-of-distribution protein–ligand complexes and identifies cryptic and allosteric pockets from amino-acid sequence alone. In parallel, Yann LeCun's critique of autoregressive models and his bet on large world models (LWMs) built on joint embedding predictive architectures (JEPA) —channelled since 2026 through AMI Labs— offers the conceptual framework needed to overcome the "snapshot" limitation of current structural predictors. In this Perspective I argue that, in HER2-positive breast cancer —where resistance to antibody–drug conjugates such as trastuzumab deruxtecan, to tyrosine-kinase inhibitors and to bispecific antibodies is governed by conformational dynamics and clonal evolution—, the convergence between the reasoning of large language models (LLMs) and the physical-chemical simulation capacity of LWMs could, between 2027 and 2030, enable the design of a new generation of therapeutic agents. I explicitly distinguish between today's demonstrated capabilities, ongoing developments and prospective projections, and I propose a roadmap for Latin American academic clinical oncology to take part in this transformation as a co-developer rather than a late adopter.

Keywords

Artificial intelligence, Drug discovery, HER2-positive breast cancer, AlphaFold, Isomorphic Labs, World models, JEPA, Antibody–drug conjugates, Computational oncology, Therapeutic resistance.

Introduction

For half a century, oncology drug discovery has been a slow, costly process dominated by a persistent attrition rate above 95%. Most of that loss comes not from medicinal chemistry but from our inability to anticipate how a molecule will behave inside a living system: how its target folds, how its binding pocket breathes, which escape pathways the tumour will activate. The underlying question is not chemical but physical and causal.

The advent of structural artificial intelligence promised to change that equation. AlphaFold solved the folding problem with near-experimental accuracy and, in doing so, opened a new era. Yet it is worth being precise about exactly what it solved: AlphaFold predicts, with remarkable fidelity, a snapshot of a protein's ground state. The biology most relevant to therapeutics —allostery, induced fit, the transient opening of cryptic pockets, the dimerisation dynamics of receptors— lives in motion, not in the still image [1].

In this Perspective I bring together three developments that, read jointly, point to the next frontier. First, the maturation of the DeepMind–Isomorphic Labs ecosystem, which has moved from predicting structures to designing ligands. Second, Yann LeCun's critique of autoregressive models and his proposal of world models as the path toward an intelligence capable of reasoning and planning. Third, the clinical setting of HER2-positive breast cancer, where therapeutic resistance poses precisely the kind of dynamic, causal problem that a molecular world model would be positioned to address. My thesis is that the integration of both worlds —the reasoning of large language models and the physical-space simulation of large world models— defines the coming decade of rational drug design in oncology.

From protein folding to rational design: the AlphaFold–Isomorphic ecosystem

AlphaFold 2 established that folding could be predicted from sequence with accuracy comparable to crystallography. AlphaFold 3 extended the scope to complexes of proteins with nucleic acids, ions and small-molecule ligands within a unified framework, reaching around 76% success on the PoseBusters benchmark. Its persistent limitation, however, is twofold: it generalises poorly to protein–ligand systems far from its training set —precisely where therapeutic innovation is decided— and it delivers static structures, blind to conformational fluctuation [2].

The Isomorphic Drug Design Engine (IsoDDE), unveiled by Isomorphic Labs in February 2026, represents a qualitative rather than merely incremental leap. Three achievements define it. In the most demanding category of the independent Runs N' Poses benchmark —systems most dissimilar to the training set— IsoDDE more than doubles the accuracy of AlphaFold 3. In binding-affinity prediction it surpasses not only earlier deep-learning methods but also gold-standard physics-based free-energy perturbation (FEP+) calculations, and it does so without requiring a crystallographic structure as input. Finally, it identifies novel pockets —orthosteric, allosteric and cryptic— from sequence alone: the paradigmatic case is the computational recapitulation of an experimentally described novel cryptic allosteric site in cereblon, with no ligand information [3].

This last point deserves interpretive emphasis. That a system trained on static structures should predict the opening of a pocket existing only in a transient conformation suggests that its internal representations have learned, implicitly, something akin to a model of the system's dynamics. I shall return to this idea, because it is the conceptual bridge of this Perspective.

The remainder of the DeepMind ecosystem completes a genuine computational-oncology infrastructure. AlphaMissense classifies close to 89% of the roughly 71 million possible missense variants of the human proteome, a tool directly applicable to the reclassification of variants of uncertain significance in genes such as BRCA1/2, PALB2, ATM and POLE. AlphaGenome carries prediction into the non-coding genome, accepting up to one million bases and predicting multiple functional modalities at single-base resolution [4,5].

The capital committed confirms the scale of the wager: a US$2.1 billion Series B in May 2026 and collaborations with Eli Lilly, Novartis and Johnson & Johnson whose stated aggregate deal value approaches US$3 billion. The company has indicated that its first candidates, arising from internal programmes in oncology and immunology, could enter clinical studies before the close of 2026. Other platforms —the open-source Boltz-2, able to predict structure and affinity roughly a thousand times faster than FEP; Generate:Biomedicines' Chroma, for de novo protein design; or Insilico Medicine's rentosertib, the first fully AI-discovered and AI-designed asset with a positive phase IIa readout— confirm that this is not an isolated phenomenon but a paradigm shift [6,7].

LeCun's critique and world models: JEPA, AMI and LWMs

While structural AI advanced, one of the founding voices of deep learning was formulating a critique that proves unexpectedly relevant to biology. Yann LeCun has consistently argued that autoregressive language models, however powerful, do not build a model of the world: they predict the next plausible token but —in his metaphor— memorising every cookbook tells you nothing about how food tastes. To reason and plan genuinely, he contends, a system needs an internal model of the physics, causality and consequences of its environment, one updated through observation and interaction [8].

His architectural proposal is the JEPA (joint embedding predictive architecture). Rather than predicting every pixel or token of the future, a JEPA learns, in a self-supervised manner, latent representations that predict future states in an abstract embedding space, where irrelevant detail is discarded and semantic and physical invariants are preserved. Formally it is trained as an energy-based model, assigning low energy to compatible continuations and high energy to incompatible ones. The image (I-JEPA) and video (V-JEPA and V-JEPA 2) variants have demonstrated a striking result: trained on unlabelled video, V-JEPA 2 learned a physical model sufficient to plan robotic actions zero-shot in novel environments, with success rates far above prior approaches [9].

In November 2025 LeCun confirmed his departure from Meta after more than a decade, and in early 2026 the founding of AMI Labs (Advanced Machine Intelligence Labs), based in Paris with initial funding of around US$1 billion. Its stated goal is to build world models —systems that learn the rules of the physical environment from multimodal sensory data— on the premise that the next advances will not come from scaling language models. A large world model, strictly speaking, learns two things in a latent space: the dynamics of the environment (how state changes as a function of actions) and a predictive model able to simulate future states across multiple timescales; inference then becomes an optimisation search in that space —that is, planning [10].

From the physical world model to the molecular world

Here lies the central argument of this Perspective. If we replace the frames of the physical world with the conformational states of a protein or a biomolecular complex, the world-model framework translates almost naturally into the molecular domain.

The "world" is the conformational and energetic landscape of the biomolecular system: the free-energy surface, its attractor basins, its collective modes of motion.

The "state" is not the Cartesian coordinate of every atom but a latent embedding capturing symmetries, slow modes and functional properties —allosteric sites, the opening of cryptic pockets, phosphorylation states.

The "action" is the introduction of a ligand, a point mutation, a post-translational modification or a change of binding partner.

"Planning" is the search, in latent space, for molecules or perturbations that drive the system toward desired states —for example, a HER2 stabilised in a closed conformation incapable of dimerising.

This reading is not entirely theoretical. IsoDDE's ability to predict induced fit and the opening of cryptic pockets outside its training set suggests the emergence of an implicit world model within its learned representations. What does not yet exist is an explicit molecular world model —trained in the manner of V-JEPA 2 on molecular-dynamics trajectories, dynamic cryo-electron microscopy and systematic experimental perturbations— able to reason about physical-chemical causality rather than mere sequence correlation. That, in my view, is the natural next step, and it is where LeCun's critique ceases to be a dispute between schools of AI and becomes a concrete biomedical research agenda.

HER2-positive breast cancer: the clinical problem and the computational challenge

Roughly 14% of female breast cancers overexpress HER2. The therapeutic arsenal has expanded extraordinarily: the monoclonal antibodies trastuzumab and pertuzumab; the antibody–drug conjugates trastuzumab deruxtecan (T-DXd) and trastuzumab emtansine; the HER2-selective tyrosine-kinase inhibitor tucatinib; the biparatopic bispecific zanidatamab; and, in the realm of targeted degraders, the recent approval of the first therapeutic PROTAC, vepdegestrant, in oestrogen-receptor-positive disease with ESR1 mutation, which opens an entirely new pharmacological modality from a regulatory standpoint [11,12].

In 2025 and 2026 the field reconfigured at an accelerating pace. The combination of T-DXd with pertuzumab was positioned in the first-line metastatic setting on the basis of the DESTINY-Breast09 trial. In early-stage disease, two recent approvals broadened the use of T-DXd: as a neoadjuvant regimen followed by trastuzumab–pertuzumab in stages II and III, and as adjuvant treatment of residual disease after neoadjuvant therapy, the latter with a roughly 53% reduction in the risk of invasive recurrence versus trastuzumab emtansine. Tucatinib retains its place after demonstrating a 34% reduction in the risk of death, with particular value in brain metastases. Zanidatamab, combined with docetaxel in the first line, achieved objective response rates above 90% [13,14].

And yet resistance remains the rule. Its mechanisms are, precisely, dynamic and causal: loss of HER2 expression or binding as escape from T-DXd; overexpression of efflux transporters such as ABCC1, an independent predictor of shorter survival; mutations in HER2's dimerisation domains —G309A, S310Y, P523S— that destabilise the interface, favour HER2:HER3 heterodimerisation and PI3K/AKT activation, and that retain sensitivity to tucatinib while losing it to other agents; activation of alternative pathways; intratumoural heterogeneity with coexisting HER2-low and HER2 3+ clones; and immune evasion. Each of these phenomena requires reasoning about how the system moves and evolves, not about a fixed image. It is, in other words, a world-model problem.

LLM + LWM convergence: applications by time horizon

Productive integration combines two complementary capabilities. Large language models contribute reasoning over the literature, hypothesis generation, integration of dispersed biomedical knowledge and the orchestration of workflows through agents. Large world models contribute the simulation of physical-chemical and conformational space, binding dynamics and planning in latent space. I propose ordering the applications by their horizon of maturity.

Already-demonstrated capabilities (now to twelve months)

The de novo design of anti-HER2 antibodies through generative AI is already a reality: de novo variants with nanomolar affinities and potency equivalent or superior to trastuzumab have been experimentally validated. Combined with IsoDDE's improvement on antibody–antigen interfaces, the design of bispecifics directed at alternative epitopes —useful against resistant variants— is computationally feasible today. The in silico reclassification of unusual HER2 mutations with AlphaMissense, and the optimisation of the drug-to-antibody ratio, linkers and payloads of antibody–drug conjugates through rapid affinity prediction, complete the immediate repertoire [15-17].

Near frontier (twelve to thirty-six months)

Three objectives are particularly attractive. First, the rational design of a second-generation antibody–drug conjugate whose payload evades recognition by ABCC1 without losing cytotoxicity: a system pairing language agents reasoning over efflux-transporter biology with a molecular model simulating the payload in the ABCC1 pocket could propose targeted modifications and shorten that cycle from years to months. Second, the design of a HER2-directed degrader: with the regulatory precedent of the first approved PROTAC, the bottleneck shifts to predicting the dynamic structure of the ternary complex —a problem that static predictors solve poorly but that IsoDDE, in light of its cereblon

performance, is beginning to address. Third, antibodies directed specifically against the mutant conformations of the dimerisation domains.

Grounded prospective vision (beyond thirty-six months)

The larger horizon is that of closed loops articulating reasoning, simulation and validation. A language-model orchestrator reasons over the literature, formulates hypotheses and plans experiments; queries a molecular world model to evaluate designs in silico; and dispatches candidates to autonomous experimental validation in automated laboratories, whose results feed back into the model. On such an architecture, simultaneous multiparametric optimisation becomes tractable —selectivity against EGFR/HER3/HER4, blood–brain-barrier penetration, conjugate stability, immunogenic profile and pharmacokinetics— and, above all, the prediction of adaptive resistance: a world model trained on clonal-evolution trajectories could anticipate which escape mutations will emerge under a given drug and design in advance the combinations that close those pathways. It would be the molecular equivalent of the zero-shot planning that world models already exhibit in the physical domain.

Limitations and challenges

Enthusiasm should be tempered with rigour. The platforms described are, to date, demonstrations of predictive capability, not of clinical efficacy; attrition in oncology has historically exceeded 95% and no computational projection suspends it by decree. The variant-prediction models were trained predominantly on data of European ancestry, which obliges us to validate their performance in Latin American populations before any clinical use —a caveat especially relevant to our practice. The regulatory framework has not yet faced dossiers in which the rationale for a molecule resides in the representations of a neural network rather than in human-interpretable pharmacology, although the recent approval of the first PROTAC suggests flexibility toward novel modalities. Equitable access is an ethical imperative: if AI-based pipelines compress timelines and costs, those benefits must reach the health systems of middle-income countries, where the burden of breast cancer is growing fastest.

Finally, it is essential to distinguish between demonstrated capabilities —IsoDDE on cereblon, V-JEPA 2 in robotics, the de novo design of anti-HER2 antibodies, the first PROTAC in the clinic— and prospective projections about how convergence will transform HER2-positive disease in the coming years. The latter is reasonable and informed, but conjectural. No experimental work has yet demonstrated an explicit molecular world model trained in the manner of V-JEPA 2; to claim otherwise would be to confuse the promise with the result. And no system, however sophisticated, eliminates the risk of hallucination: the physician must remain the ultimate reviewer of every hypothesis these tools generate [18,19].

Conclusion

Oncology stands at the convergence of two lineages of artificial intelligence that until recently ran in parallel. The first, embodied in AlphaFold and IsoDDE, taught machines to see the structure of life. The second, articulated in LeCun's critique and the JEPA architectures, seeks to teach them to understand its dynamics and to plan upon it. HER2-positive breast cancer, with its rich pharmacology and its tenacious resistance, is an ideal proving ground for that synthesis. The task of the academic clinical community —and most especially the Latin American one— is not to await the result but to generate the data, train the people and build the consortia that will let us take part in this transformation as protagonists. We already have the snapshot; what comes next is learning to read the film.

Declaration on the use of artificial intelligence

In preparing this manuscript, generative artificial-intelligence tools based on large language models were used as assistants for the search and synthesis of scientific literature, the conceptual organisation of the arguments and the editing of the text. The author conceived the central thesis, defined the scope, verified and interpreted the sources, and critically reviewed every section. The author assumes full responsibility for the content, claims and conclusions expressed herein. No artificial-intelligence tool is listed as an author or credited with authorship, in line with the current recommendations of the International Committee of Medical Journal Editors (ICMJE). The author has checked the cited references against their primary sources and encourages readers to consult the original articles for methodological detail.

References

1. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021; 596: 583-589.
2. Abramson J, Adler J, Dunger J, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024; 630: 493-500.
3. Isomorphic L. The Isomorphic Labs Drug Design Engine unlocks a new frontier beyond AlphaFold. Technical report. London: Isomorphic Labs. 2026.
4. Cheng J, Novati G, Pan J, et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. Science. 2023; 381: eadg7492.
5. Avsec Ž, Latysheva N, Cheng J, et al. A unifying model of the regulatory genome with AlphaGenome. Nature. 2026; 649: 1206-1218.
6. Passaro S, Corso G, Wohlwend J, et al. Boltz-2: towards accurate and efficient binding affinity prediction. bioRxiv. 2025.
7. Ingraham JB, Baranov M, Costello Z, et al. Illuminating protein space with a programmable generative model. Nature. 2023; 623: 1070-1078.
8. LeCun Y. A path towards autonomous machine intelligence. OpenReview; 2022. Version 0.9.2.
9. Assran M, Bardes A, Fan D, et al. V-JEPA 2: self-supervised video models enable understanding, prediction and planning. Meta AI; 2025.

10. Balestriero R, LeCun Y. LeJEPA: provable and scalable self-supervised learning without the heuristics. 2025.
11. Loibl S, Mano M, Untch M, et al. Trastuzumab deruxtecan in HER2-positive early breast cancer with residual invasive disease after neoadjuvant therapy (DESTINY-Breast05). N Engl J Med. 2025.
12. Murthy RK, Loi S, Okines A, et al. Tucatinib, trastuzumab, and capecitabine for HER2-positive metastatic breast cancer (HER2CLIMB). N Engl J Med. 2020;3 82: 597-609.
13. Wang X, Liu Y, Sun T, et al. Zanidatamab in combination with docetaxel in first-line HER2-positive breast cancer: an open-label, multicentre, phase Ib/II study. ESMO Open. 2025; 10: 105852.
14. Chen W, Hata AN, Juric D, et al. Trastuzumab deruxtecan resistance via loss of HER2 expression and binding. Cancer Discov. 2026; 16: 235-249.
15. Shanehsazzadeh A, Bachas S, McPartlon M, et al. Unlocking de novo antibody design with generative artificial intelligence. bioRxiv. 2023.
16. US Food and Drug Administration. FDA approves vepdegestrant for ER-positive, HER2-negative advanced breast cancer with ESR1 mutations. Silver Spring. 2026.
17. US Food and Drug Administration. FDA approves two separate indications for fam-trastuzumab deruxtecan-nxki in HER2-positive early-stage breast cancer. Silver Spring. 2026.
18. Tommasi C, Airò Farulla L, Curigliano G, et al. Overcoming resistance and optimizing sequential anti-HER2 therapy in HER2-positive breast cancer. J Clin Med. 2024; 13: 1873. DOI:10.3390/jcm13071873.
19. Sun Q, Melnychuk H, et al. Costs and causes of oncology drug attrition. JAMA Netw Open. 2023. PMCID: PMC10383012.