Research Article

International Journal of Tumor Research

From Protein Folding to Rational Drug Design: The DeepMind–Isomorphic Labs Ecosystem and the Dawn of Computational Oncology

Enrique Díaz Cantón, MD, MSc, MSc[1,2*]

[1]*Department of Medical Oncology and Artificial Intelligence in Medicine, School of Medicine, CEMIC University Institute, Buenos Aires, Argentina.*

[2]*Postgraduate Program in AI and Medicine, Academia Nacional de Medicina, Buenos Aires, Argentina.*

[*]Correspondence:
Enrique Díaz Cantón, MD, MSc, MSc, Department of Medical Oncology and Artificial Intelligence in Medicine, School of Medicine, CEMIC University Institute, Buenos Aires, Argentina

Received: 13 Jan 2026; **Accepted:** 18 Feb 2026; **Published:** 04 Mar 2026

Citation: Enrique Díaz Cantón. From Protein Folding to Rational Drug Design: The DeepMind–Isomorphic Labs Ecosystem and the Dawn of Computational Oncology. Int J Tumor Res. 2026; 2(1): 1-6.

ABSTRACT

On February 10, 2026, Isomorphic Labs unveiled the Isomorphic Labs Drug Design Engine (IsoDDE), a unified computational system that more than doubles the accuracy of AlphaFold 3 in protein-ligand structure prediction for novel systems and surpasses gold-standard physics-based methods in binding affinity estimation. This milestone does not emerge in isolation: it represents the culmination of a vertically integrated artificial intelligence ecosystem built by Google DeepMind and Isomorphic Labs—spanning AlphaFold for structural biology, AlphaMissense for coding variant pathogenicity, and AlphaGenome for non-coding regulatory variant interpretation—that collectively promises to transform how we discover and develop cancer therapeutics. This editorial examines IsoDDE's capabilities in the context of this ecosystem, argues that its demonstrated capacity to model induced fit and cryptic pocket dynamics suggests the emergence of implicit world models in molecular prediction, and discusses the profound implications for oncology: from reclassifying variants of uncertain significance to rationally designing drugs against historically undruggable targets. While significant translational challenges remain, the convergence of these technologies heralds a paradigm shift from empirical to computational-first drug discovery in oncology.

Keywords
Artificial intelligence, Isomorphic Labs, AlphaFold, AlphaMissense, AlphaGenome, Drug design engine, Oncology, Precision medicine, World models, Computational drug discovery.

A New Chapter in the AI–Biology Revolution

The 2024 Nobel Prize in Chemistry, awarded to Demis Hassabis and John Jumper for AlphaFold, marked the formal recognition that artificial intelligence had solved one of biology's grand challenges: the protein folding problem [1]. But for those of us in clinical oncology, that recognition carried a more urgent subtext. Cancer drug development remains one of medicine's most vexing bottlenecks: a decade-long, multi-billion-dollar gauntlet with approximately 95% attrition rates in clinical trials. If AI could predict how proteins fold, could it also predict how drugs bind—and ultimately, how to design them rationally?

On February 10, 2026, Isomorphic Labs provided a compelling answer. The Isomorphic Labs Drug Design Engine (IsoDDE) is not merely an incremental improvement over AlphaFold 3; it represents a qualitative leap into a new paradigm of predictive drug design [2]. By more than doubling the structure prediction accuracy of AlphaFold 3 on the most challenging novel systems, surpassing gold-standard physics-based methods in binding affinity prediction, and demonstrating the ability to identify cryptic and allosteric binding pockets from amino acid sequence alone, IsoDDE bridges the gap that has separated computational structure prediction from actionable drug discovery.

Critically, IsoDDE does not exist in a vacuum. It is the apex of a vertically integrated AI ecosystem that Google DeepMind and Isomorphic Labs have systematically assembled over five years: AlphaFold for protein structure prediction [1,3], AlphaMissense for coding variant pathogenicity classification [4], and AlphaGenome

for non-coding regulatory variant effect prediction [5]. Together, these tools constitute what I would argue is the most consequential technological platform for biomedicine since the completion of the Human Genome Project. And for oncology specifically, the implications are transformative.

AlphaFold: The Foundation That Changed Everything

AlphaFold 2 solved a 50-year-old grand challenge by predicting protein three-dimensional structures with near-experimental accuracy from amino acid sequence alone [1]. Its successor, AlphaFold 3, expanded this capability to predict the structures of protein complexes with DNA, RNA, small molecule ligands, and ions within a unified framework [3]. The AlphaFold Protein Structure Database has been used by over 3 million researchers across more than 190 countries, and approximately 30% of related research focuses specifically on disease biology.

For oncology, AlphaFold's impact was immediate and profound. The CDK20 case study—where researchers from DeepMind and Insilico Medicine used AlphaFold to elucidate the structure of a liver cancer target and generate a lead inhibitor within 30 days [6]—demonstrated that timelines traditionally measured in years could be compressed to weeks. AlphaFold 3's ability to model protein-ligand interactions with a 76% success rate on the PoseBusters benchmark opened the door to computational screening of billions of candidate compounds against cancer targets.

However, critical limitations remained. Subsequent benchmarks revealed that AlphaFold 3 struggled to generalise to protein-ligand systems dissimilar to its training data—precisely the unexplored regions of biomolecular space where the biggest opportunities in drug discovery reside [2,7]. The model predicted static, ground-state structures, but real biological function depends on dynamics: proteins breathe, flex, and undergo induced conformational changes upon ligand binding. For oncology drug design, where targets like KRAS or p53 have historically been considered "undruggable" precisely because of their dynamic conformational landscapes, this was more than an academic limitation—it was a translational barrier.

AlphaMissense: Decoding the Language of Oncogenic Variation

While AlphaFold addressed the structural dimension, precision oncology has been confronted with an equally formidable interpretive challenge: the vast majority of genetic variants encountered in clinical tumor sequencing remain classified as variants of uncertain significance (VUS). Of approximately 71 million possible missense variants in the human proteome, only 0.1% had been definitively classified by human experts prior to AlphaMissense [4].

AlphaMissense, published in Science in 2023, adapted the AlphaFold architecture to predict variant pathogenicity, achieving 89% classification of all possible human missense variants as either likely pathogenic or likely benign [4]. For oncology, this has direct clinical relevance. In DNA damage repair (DDR) genes—ATM, BRCA1, BRCA2, PALB2, POLE—VUS reclassification can determine whether a patient is eligible for PARP inhibitor therapy or qualifies for enhanced screening protocols. AlphaMissense scores for DDR gene variants presented at ASCO 2024 demonstrated remarkable concordance with expert-curated clinical classifications, offering a scalable solution to a problem that tumor boards face daily.

Yet AlphaMissense, by design, addresses only the 2% of the genome that encodes proteins. The remaining 98%—the non-coding regulatory genome, once dismissively labeled "junk DNA"—harbors the majority of disease-associated genetic variation identified by genome-wide association studies. For oncology, where enhancer hijacking, promoter mutations, and aberrant splicing increasingly emerge as cancer drivers, a tool capable of interpreting non-coding variants was urgently needed.

AlphaGenome: Illuminating the Dark Genome of Cancer

Published in Nature in January 2026, AlphaGenome represents the logical and necessary extension of this ecosystem into the non-coding regulatory genome [5]. The model accepts up to one million base pairs of DNA sequence as input and predicts thousands of functional genomic readouts at single-base-pair resolution across 11 modalities: gene expression, transcription initiation, chromatin accessibility, histone modifications, transcription factor binding, chromatin contact maps, splice site usage, and splice junction coordinates [5]. In benchmarking, AlphaGenome matched or exceeded the strongest available external models in 25 of 26 variant effect prediction tasks.

The oncological implications are immediate and far-reaching. AlphaGenome can score the impact of a single variant across thousands of molecular features in approximately one second [5,8]. As demonstrated in the original publication, the model accurately recapitulated the mechanisms of clinically relevant variants near the TAL1 oncogene—a locus frequently implicated in T-cell acute lymphoblastic leukemia through non-coding regulatory mutations [5]. Early adopters have already used AlphaGenome to pinpoint mutations in cancer genomes that drive proliferation, and nearly 3,000 scientists across 160 countries submitted approximately 1 million daily requests during the initial release period [8,9].

AlphaGenome fills the critical gap between AlphaMissense's protein-coding variant interpretation and the broader genomic landscape that clinical oncologists increasingly need to navigate. Together, these tools create a comprehensive AI-powered variant interpretation pipeline spanning the entire genome—from coding mutations to regulatory non-coding variants—that could fundamentally reshape how we annotate and prioritize genetic alterations in tumor sequencing reports.

The Isomorphic Labs Drug Design Engine: Where Understanding Becomes Design

If AlphaFold, AlphaMissense, and AlphaGenome constitute the interpretive arm of this ecosystem, IsoDDE represents its

translational engine—the point where computational understanding is converted into rational drug design [2].

The technical report reveals three capabilities that merit particular attention from the oncology community [2]. First, on the Runs N' Poses benchmark—designed to test generalisation to truly novel protein-ligand systems—IsoDDE more than doubles AlphaFold 3's accuracy on the hardest generalisation category, achieving a 50% success rate on systems most dissimilar to the training set where AlphaFold 3 largely fails. This is not an incremental improvement; it is a step change in the ability to computationally model first-in-class drug targets.

Second, IsoDDE's binding affinity predictions surpass all deep-learning methods and, remarkably, exceed gold-standard free energy perturbation (FEP+) physics-based methods on curated benchmarks—without requiring experimental crystal structures as input [2]. In traditional drug development, FEP calculations are computationally expensive, require significant manual setup for each system, and remain accessible primarily to well-resourced pharmaceutical companies. IsoDDE democratises this capability, delivering experimental-grade precision at a fraction of the time and cost.

Third, and perhaps most consequential for oncology, IsoDDE demonstrates the ability to identify novel ligandable pockets from amino acid sequence alone—including cryptic pockets that are hidden in the unliganded protein state [2]. The cereblon example is illustrative: for 15 years, a single thalidomide-binding pocket was believed to be the principal druggable site on this E3 ligase component central to molecular glue degrader therapy. A 2026 experimental study discovered a novel allosteric cryptic pocket [10]. IsoDDE recapitulated this discovery computationally, predicting both the known and novel sites using only the protein sequence as input, without specifying ligand identity [2]. This capacity to computationally reveal druggable pockets that eluded 15 years of experimental investigation has obvious implications for historically intractable oncology targets.

Furthermore, the antibody-antigen structure prediction capabilities of IsoDDE—outperforming AlphaFold 3 by 2.3-fold and Boltz-2 by 19.8-fold in high-fidelity predictions—open transformative possibilities for computational de novo design of therapeutic antibodies, bispecifics, and antibody-drug conjugates, therapeutic modalities that now constitute one of the fastest-growing segments of the oncology pipeline [2].

Beyond Static Prediction: The Implicit Emergence of World Models

There is a deeper conceptual significance to IsoDDE's achievements that deserves explicit articulation. The successful prediction of induced fit phenomena—where a protein substantially changes its three-dimensional conformation to accommodate a bound ligand—and the identification of cryptic pockets that only exist in the presence of a binding partner represent something more profound than improved benchmark scores. These capabilities suggest the emergence of what might be termed an implicit world model within the system's learned representations.

I have argued previously that the principal limitation of AlphaFold—predicting a static ground-state "photograph" of a protein rather than its dynamic conformational landscape—represents a fundamental blind spot for drug design [11]. The induced fit problem is not merely geometric; it is physical and thermodynamic. A drug does not simply dock into a pre-existing pocket; it engages in a complex energetic dialogue with the protein, triggering conformational rearrangements that propagate through the molecular structure. As noted in the AlphaMissense paper itself, AlphaFold 2 is "largely insensitive to the variation of the input sequence and cannot accurately predict structural changes following a point mutation" [4].

Yann LeCun's Joint Embedding Predictive Architecture (JEPA) framework provides a theoretical lens through which to interpret IsoDDE's advances [12]. World models, in LeCun's formulation, learn abstract latent representations that capture the underlying physics and causal structure of a system, enabling prediction of future states without exhaustive simulation of every molecular degree of freedom. When IsoDDE correctly predicts that a ligand will open a cryptic pocket that is invisible in the unliganded structure—an event outside its training set—it is, in effect, performing a form of latent physical reasoning about conformational dynamics.

This observation has profound implications for oncology drug design. Many of the most clinically important cancer targets—mutant KRAS, wild-type p53 stabilizers, MYC inhibitors—have been classified as "undruggable" precisely because they lack stable, well-defined binding pockets in their ground-state conformations. If IsoDDE has indeed learned an implicit world model capable of reasoning about conformational dynamics, the boundary of what is "druggable" may be far more expansive than the field has assumed. The explicit development of molecular world models—AI systems that predict not just what structure a protein adopts, but how and why conformational transitions occur—represents, in my view, the next critical frontier for AI in drug discovery.

What This Means for Oncology: A Paradigm Shift in Five Dimensions

The convergence of AlphaFold, AlphaMissense, AlphaGenome, and IsoDDE creates a paradigm shift for oncology that operates across at least five dimensions.

The first is target discovery. IsoDDE's cryptic pocket identification capability, combined with AlphaGenome's ability to prioritize non-coding regulatory drivers of cancer, means that the space of potentially actionable cancer targets expands dramatically. Allosteric and cryptic binding sites on proteins previously considered undruggable become computationally accessible. The cereblon example—where IsoDDE predicted a pocket that eluded 15 years of experimental investigation—is a proof of principle with immediate relevance for molecular glue degrader therapy,

one of the most promising emerging modalities in oncology [2,10].

The second dimension is variant interpretation. The combination of AlphaMissense for coding variants and AlphaGenome for non-coding regulatory variants creates, for the first time, a genome-wide computational interpretation layer for clinical tumor sequencing. The era in which 40–60% of variants in a next-generation sequencing report are classified as VUS—leaving clinicians and patients in an uncomfortable limbo of uncertainty—may be approaching its end.

Third, IsoDDE's ability to predict binding affinities with experimental-grade accuracy at computational speed enables rational drug optimisation across vast chemical and biological spaces. For antibody-based therapies—including antibody-drug conjugates (ADCs), bispecific T-cell engagers (BiTEs), and immune checkpoint modulators—the 2.3-fold improvement over AlphaFold 3 in antibody-antigen interface prediction opens the door to de novo computational design rather than the empirical screen-and-iterate approach that currently dominates.

Fourth, timeline compression. The traditional drug discovery paradigm—from target identification through lead optimization to candidate selection—typically requires 4–6 years. IsoDDE, combined with the partnership framework that Isomorphic Labs has established with Eli Lilly and Novartis (valued at nearly $3 billion), is designed to collapse these timelines to months [13]. The report that Isomorphic Labs is preparing AI-designed molecules for clinical evaluation represents the first test of whether computational drug design can deliver on this promise.

Fifth, and most fundamentally, the epistemological shift. These tools collectively transition oncology from a primarily empirical science—where drug discovery depends on high-throughput screening, structure-activity relationship intuition, and iterative experimental cycles—toward a predictive, computational-first discipline. We are witnessing the emergence of what might be called "in silico oncology": a practice in which computational prediction precedes and guides experimental validation, rather than the reverse.

Challenges and the Road Ahead

Intellectual honesty demands that we temper enthusiasm with appropriate caveats. IsoDDE's technical report, while impressive, describes predictive capabilities—not clinical outcomes. The ultimate validation of any drug design engine is whether the molecules it designs prove safe and efficacious in human patients. The attrition rate in oncology drug development is notoriously high, and computational predictions, however accurate, cannot fully capture the complexity of in vivo pharmacology, toxicology, and tumor heterogeneity.

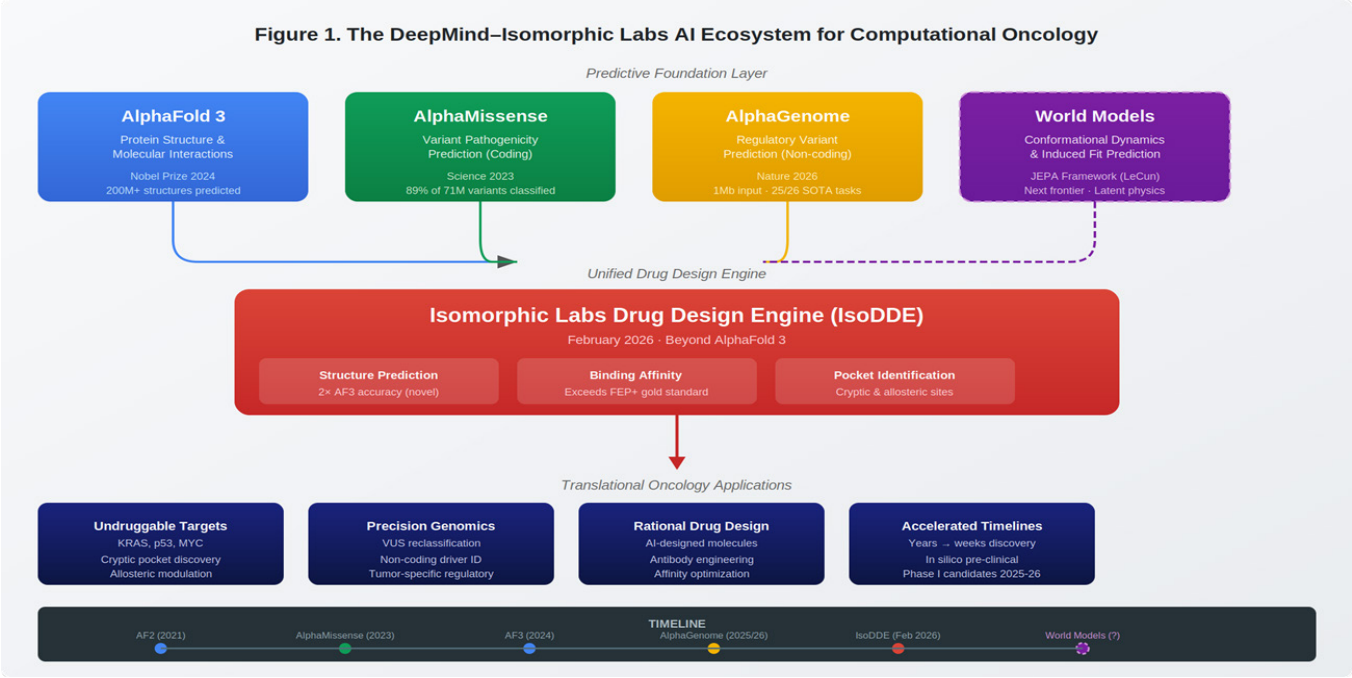

Figure 1: The DeepMind–Isomorphic Labs AI ecosystem for computational oncology. The predictive foundation layer comprises AlphaFold 3 (protein structure and molecular interactions), AlphaMissense (coding variant pathogenicity), AlphaGenome (non-coding regulatory variant effects), and the emerging frontier of molecular World Models for conformational dynamics prediction. These foundations converge in the Isomorphic Labs Drug Design Engine (IsoDDE), which integrates structure prediction, binding affinity estimation, and pocket identification into a unified drug design platform. The translational oncology applications span four domains: undruggable target engagement through cryptic pocket discovery, precision genomics through comprehensive variant interpretation, rational drug design for small molecules and biologics, and accelerated development timelines. The dashed border around World Models indicates this as a proposed next frontier.

AlphaGenome's predictions, while state-of-the-art, require careful validation across diverse populations and disease contexts before clinical implementation [5,8]. The model was trained primarily on European-ancestry data, and performance in underrepresented populations must be rigorously assessed. Similarly, AlphaMissense scores, while highly concordant with expert classifications, should be integrated into, not substituted for, clinical interpretation frameworks that consider the full biological and clinical context.

Regulatory frameworks are also evolving to accommodate AI-designed therapeutics, but have not yet been stress-tested with fully computationally designed molecules. How regulatory agencies will evaluate molecules where the "rationale" is embedded in neural network representations rather than human-interpretable pharmacological reasoning remains an open and consequential question.

Finally, equity concerns must be addressed head-on. If AI-driven drug design dramatically compresses timelines and reduces costs, these benefits must reach patients in low- and middle-income countries where cancer burden is growing most rapidly. The decision by Google DeepMind to make AlphaFold, AlphaMissense, and AlphaGenome freely available for non-commercial use is laudable and sets an important precedent. However, Isomorphic Labs' commercial drug design capabilities remain proprietary, and how the fruits of this technology will be distributed globally is a question that the oncology community must actively engage with.

Conclusion: The Beginning of Computational Oncology

We are standing at what future historians of medicine may recognize as a pivotal inflection point. The publication of IsoDDE, arriving just weeks after AlphaGenome's formal publication in Nature, completes a technological ecosystem that spans from DNA sequence interpretation through protein structure prediction to rational drug design. For the first time, a unified computational pipeline can read a cancer patient's genomic alterations—coding and non-coding—predict their functional consequences at the protein and regulatory level, identify druggable pockets on the resulting aberrant proteins—including cryptic sites invisible to experimental methods—and rationally design molecules to engage those targets with predicted binding affinities that rival experimental measurement.

This does not mean that experimental biology, clinical trials, or human expertise will become obsolete. Rather, the relationship between computation and experimentation is being inverted: AI generates hypotheses and designs candidates, while experiments validate and refine. The next frontier—the development of explicit molecular world models that predict not just static endpoints but the causal physics of conformational dynamics—will further expand what is computationally tractable.

As oncologists, we have spent decades navigating the limitations of empirical drug discovery: the undruggable targets, the variants of uncertain significance, the years-long development timelines, the crushing attrition rates. The DeepMind–Isomorphic Labs ecosystem does not eliminate these challenges overnight, but it provides, for the first time, a plausible computational path through them. The age of computational oncology is not coming. It has arrived.

Declaration of Artificial Intelligence Use

This manuscript was prepared with the assistance of Claude (Anthropic, Claude Opus 4 model), used as a support tool for literature search, content organization, translation, and drafting of the initial version. The intellectual content, editorial opinions, conceptual framework, scientific arguments, and final critical revision are entirely the responsibility of the author. The factual information and bibliographic references were independently verified by the author. The use of generative AI is disclosed in accordance with current recommendations on transparency in scientific publications.

References

1. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021; 596: 583-589. doi:10.1038/s41586-021-03819-2
2. Isomorphic Labs Team. Accurate predictions of novel biomolecular interactions with IsoDDE. Technical Report. February 10, 2026. Available at: https://storage.googleapis.com/isomorphiclabs-website-public-artifacts/isodde_technical_report.pdf
3. Abramson J, Adler J, Dunger J, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. Nature. 2024; 630: 493-500. doi:10.1038/s41586-024-07487-w
4. Cheng J, Novati G, Pan J, et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. Science. 2023; 381: 7492. doi:10.1126/science.adg7492
5. Avsec Ž, Latysheva N, Cheng J, et al. Advancing regulatory variant effect prediction with AlphaGenome. Nature. 2026; 649: 1206-1218. doi:10.1038/s41586-025-10014-0
6. Ren F, Ding X, Zheng M, et al. AlphaFold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel CDK20 small molecule inhibitor. Chem Sci. 2023; 14: 1443-1452. doi:10.1039/d2sc05709c
7. Škrinjar P, Sommer K, Gligorijevic V, et al. Runs N' Poses: a benchmark for protein-ligand cofolding generalisation. bioRxiv. 2025. doi:10.1101/2025.02.03.636309
8. Molteni M. DeepMind releases AlphaGenome source code, aiming to improve its AI predictions of gene regulation. STAT News. January 28, 2026.
9. Google DeepMind. AlphaGenome: AI for better understanding the genome. Blog post. June 25, 2025. Available at: https://deepmind.google/blog/alphagenome-ai-for-better-understanding-the-genome/
10. Dippon J, et al. Discovery of a novel allosteric cryptic pocket on cereblon. Nature. 2026. doi:10.1038/s41586-025-09994-w

11. Diaz Canton E. Beyond the Syntax of the Genome: Toward Predictive World Models in Triple Negative Breast Cancer. MAR Oncology and Hematology. 2026; 6: 1-5.
12. LeCun Y. A path towards autonomous machine intelligence. Version 0.9.2. June 27, 2022. Available at: https://openreview.net/pdf?id=BZ5a1r-kVsf
13. Isomorphic Labs. Partnerships: transforming drug discovery with AI. Available at: https://www.isomorphiclabs.com/partnerships