MAR Oncology and Hematology **(**2026**)** 6:2
Editorial Article

Beyond the Syntax of the Genome: Toward Predictive World Models in Triple-Negative Breast Cancer

Enrique Díaz Cantón, MD, (MSc)*

***Correspondence to**: **Enrique Díaz Cantón**, Professor of Oncology and Artificial Intelligence in Medicine, Instituto Universitario CEMIC. Director, Postgraduate Program in AI and Medicine, Academia Nacional de Medicina, Argentina.

Received: 18 January 2026

Published: 01 February 2026

DOI: https://doi.org/10.5281/zenodo.18390955

Abstract

Current artificial intelligence applications in triple-negative breast cancer (TNBC) rely predominantly on generative architectures that synthesize existing knowledge but lack the capacity to model tumor dynamics prospectively. This editorial argues that TNBC, with its characteristic clonal plasticity and therapeutic evasion, may require a fundamentally different computational approach. Drawing on recent advances in predictive architectures—particularly joint-embedding models and world models—I propose that the future of AI in TNBC lies in systems capable of anticipating tumor trajectories through biological state space, rather than merely pattern-matching against historical data. While this remains a hypothesis requiring empirical validation, the conceptual shift from reactive to predictive modeling may prove essential for managing cancers that behave as adaptive, navigating agents.

Keywords: *Triple-negative breast cancer; Joint-embedding predictive architecture; World models; Clonal evolution; Tumor heterogeneity; Artificial intelligence.*

The Limits of the Linguistic Mirror

We have achieved remarkable success in cataloging the molecular landscape of triple-negative breast cancer. We map mutations, profile transcriptomes, and characterize the tumor microenvironment with ever-increasing resolution. Yet our therapeutic progress has not kept pace with our descriptive capabilities. TNBC remains the most lethal breast cancer subtype, accounting for a disproportionate share of breast cancer mortality despite representing only 15% of cases.[1]

Part of this gap, I suggest, stems from a mismatch between our computational tools and the nature of the disease. Current AI applications in oncology—dominated by large language models and pattern-recognition systems—excel at synthesizing what has been observed. They function as sophisticated mirrors of accumulated clinical experience: given a patient profile, they retrieve and recombine relevant precedents from the literature and trial databases.

But TNBC is not a static pattern to be recognized. It is a dynamic system that evolves under selective pressure, exploiting metabolic flexibility and phenotypic plasticity to evade therapeutic intervention.[2] As Yann LeCun has argued in a different context, generative AI—however impressive—fundamentally operates by predicting probable sequences based on training data.[3] This is not the same as understanding the causal dynamics of a system. To borrow his framing: predicting the next token in a medical report is not equivalent to modeling the "physics" of tumor progression.

TNBC as a Navigating Agent

What makes TNBC computationally distinctive is its behavior as what might be called a "navigating agent." Unlike more indolent cancers that follow relatively predictable trajectories, TNBC actively explores biological state space, seeking paths around therapeutic obstacles. Through epithelial-mesenchymal transition, metabolic reprogramming, and clonal selection, the tumor continuously repositions itself in response to environmental pressures.[4,5]

This framing suggests that effective AI for TNBC management would need to do more than recognize patterns—it would need to anticipate trajectories. The question is not merely "what does this tumor look like?" but "where is this tumor going, and how can we block its path?"

Recent work in machine learning offers a conceptual framework that may be relevant here. LeCun's Joint-Embedding Predictive Architecture (JEPA) represents a departure from generative approaches.[3,6] Rather than learning to produce outputs token-by-token, JEPA systems learn to predict the latent representation of future states from current observations. They operate in an abstract embedding space, filtering out irrelevant details to focus on the structural variables that drive system dynamics.

World Models for Tumor Dynamics: A Hypothesis

The application of world models and predictive architectures to oncology remains largely unexplored. What I propose here is therefore a hypothesis, not a demonstrated capability. But the conceptual alignment is suggestive.

A world model, in the machine learning sense, is an internal representation that allows an agent to simulate the consequences of actions before taking them.[6] Applied to TNBC, such a model would not attempt to predict every molecular fluctuation. Instead, it would learn the underlying dynamics that govern therapeutic response and resistance: which signaling dependencies matter, how the tumor microenvironment shapes evolutionary trajectories, what metabolic "detours" become available when primary pathways are blocked.

In principle, such a system could anticipate resistance mechanisms before they manifest clinically. Rather than waiting for PI3K/AKT/mTOR inhibitor failure and then characterizing the escape route, a predictive model might identify the most probable evasion pathways in advance, enabling preemptive combination strategies.

The data requirements for such an approach would be substantial. We would need longitudinal, multimodal profiling of tumors under treatment—genomic, transcriptomic, metabolomic, and imaging data collected at multiple timepoints. Current datasets rarely provide this kind of dynamic information. But as liquid biopsy technologies mature and real-world data infrastructure improves, the necessary observational basis may become available.

Scaling Clinical Intelligence

One advantage of computational approaches, as Geoffrey Hinton has observed, is their capacity for knowledge sharing.[7] Human expertise is bound to the individual: my experience managing TNBC in Buenos Aires cannot be directly transferred to a colleague in Tokyo. It must be communicated through language—a slow, lossy process that cannot capture the full texture of clinical intuition.

A predictive modeling framework, by contrast, could in principle learn from every TNBC case globally, integrating patterns across populations and treatment contexts that no individual clinician could synthesize. This is not a replacement for clinical judgment but an augmentation of it—a way of scaling the observational basis on which our decisions rest.

The Irreducible Role of the Oncologist

Even if predictive architectures fulfill their theoretical promise, the role of the oncologist will remain essential—though perhaps transformed. Algorithms may eventually model tumor trajectories with high accuracy, but the clinical encounter encompasses dimensions that no computational system captures.

The patient facing a metastatic TNBC diagnosis is navigating an existential crisis, not merely a biological one. She brings values, fears, relationships, and life goals that must inform treatment decisions. The oncologist's role becomes one of integration: synthesizing computational predictions with clinical judgment, patient preferences, and the irreducible uncertainty that accompanies serious illness.

In this framework, we are not replaced by machines but liberated by them—freed from the cognitive burden of pattern-matching against vast datasets, we can focus on the aspects of care that remain distinctly human: communication, deliberation, and the stewardship of patients through their most difficult passages.

Conclusion

The challenge of TNBC may require us to move beyond the current paradigm of AI as a sophisticated mirror of accumulated knowledge. Tumors that behave as adaptive, navigating agents may demand predictive systems capable of modeling their dynamics prospectively—anticipating where they are going rather than merely recognizing where they have been.

Joint-embedding architectures and world models offer a conceptual framework for this shift, though their application to oncology remains to be demonstrated. The path from hypothesis to clinical tool will require substantial investment in longitudinal data collection, computational infrastructure, and rigorous validation. But as we consider the future of AI in TNBC, we should ask whether incremental refinements to existing approaches will suffice—or whether the nature of the disease demands a more fundamental rethinking of how we model it.

Diclosure: The author utilized generative artificial intellingence for sintax and linguistic accurance

References

1. Bianchini G, De Angelis C, Licitra L, et al. Treatment landscape of triple-negative breast cancer—expanded options, evolving needs. Nat Rev Clin Oncol. 2022;19(2):91-113.

2. Garrido-Castro AC, Lin NU, Polyak K. Insights into molecular classifications of triple-negative breast cancer: improving patient selection for treatment. Cancer Discov. 2019;9(2):176-198.

3. LeCun Y. A path towards autonomous machine intelligence. Version 0.9.2. OpenReview. 2022. https://openreview.net/pdf?id=BZ5a1r-kVsf

4. Dongre A, Weinberg RA. New insights into the mechanisms of epithelial-mesenchymal transition and implications for cancer. Nat Rev Mol Cell Biol. 2019;20(2):69-84.

5. Faubert B, Solmonson A, DeBerardinis RJ. Metabolic reprogramming and cancer progression. Science. 2020;368(6487):eaaw5473.

6. LeCun Y. Objective-driven AI: towards machines that can reason, plan, and learn. Keynote address, World Artificial Intelligence Conference; 2024; Shanghai, China.

7. Hinton G. Will digital intelligence replace biological intelligence? Romanes Lecture, University of Oxford; February 2024.