OPEN ACCESS

Precision Medicine COMMENTARY

Redefining the tumor digital twin through cross-domain manifold mapping

Yuanning Hu,[1] Wei Zong,[2,3] Juncheng Wang,[1] Manqi Cai,[4] and Jian Zou[1,*]

[1]Department of Statistics, School of Public Health, Chongqing Medical University, Chongqing 400016, China

[2]Department of Biostatistics and Health Data Science, School of Public Health, University of Pittsburgh, PA 15260, USA

[3]School of Data Science, The Chinese University of Hong Kong, Shenzhen 518172, China

[4]C3 AI, RedwoodCity, California 94063, USA

*Correspondence: jianzou@cqmu.edu.cn

Received: February 16, 2026; Accepted: April 23, 2026; Published Online: April 24, 2026; https://doi.org/10.59717/j.xinn-oncol.2026.100012

© 2026 The Author(s). This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

Citation: Hu Y., Zong W., Wang J., et al. (2026). Redefining the tumor digital twin through cross-domain manifold mapping. *The Innovation Oncology* **1**:100012.

THE CRISIS OF TRANSLATION AND THE IMPERATIVE OF CONGRUENCE

Modern oncology is characterized by a widening chasm between scientific innovation and clinical productivity. Despite an unprecedented explosion in genomic data and a proliferation of preclinical models—ranging from cell lines, patient-derived xenografts (PDX) to organoids—the efficiency of drug development remains critically low. Approximately 90% of drug candidates entering clinical trials fail to receive approval, with oncology experiencing some of the highest attrition rates.

This translational impasse is compounded by conflicting evidence on the reliability of preclinical models themselves. We previously reconciled a long-standing debate on whether murine models faithfully mimic human biology by demonstrating that congruence is pathway-dependent: model organisms can faithfully mimic human biology in specific pathways (e.g., cell cycle regulation) while simultaneously diverging in others (e.g., specific immune responses).[1] The translational bottleneck, therefore, is not a failure of the models themselves, but a failure to rigorously quantify where and to what degree a model is congruent with human biology.

Recent computational work has formalized this finding for cancer research. Domain adaptation methods such as PRECISE have shown that drug response predictors trained on cell lines can transfer to patient tumors when aligned through shared latent subspaces.[2] Complementary frameworks (e.g. CASCAM) have further demonstrated that model selection requires resolving congruence at the level of individual biological mechanisms, not aggregate similarity.[3] Together, these studies frame the problem not as a lack of models, but as a failure to quantify which axes of biology a given model faithfully represents.

THE ILLUSION OF THE *AB INITIO* DIGITAL TWIN

To mitigate translational risk, the field is increasingly turning toward "Digital Twins"—dynamic, evolving virtual representations of a patient's malignancy that would allow clinicians to simulate drug responses *in silico* before administering toxic therapies.[4] Proposed approaches range from mechanistic computational simulators to statistical virtual patient models, but the most ambitious route is the *ab initio* construction: building a virtual cell from scratch, either through mechanistic equations that encode biochemical kinetics or through massive foundation models trained on raw molecular data.

While intellectually compelling, this approach faces two fundamental bottlenecks. First is the "parameter identifiability" problem: a single human cell involves thousands of reactions and parameters, and many are not practically identifiable from sparse patient data. Estimating these parameters for a specific individual from a simple biopsy is mathematically precarious. Second is the "Clever Hans" effect: deep learning models trained on raw biological data often learn technical artifacts (batch effects, culture-condition proxies) rather than biological signals. Because drug response can shift markedly with culture conditions, models trained on heterogeneous *in vitro* datasets risk learning culture-condition proxies rather than tumor-intrinsic drivers if those factors correlate with phenotypic labels. Furthermore, the domain shift between an *in vitro* cell line and the hypoxic, immune-active *in vivo* tumor microenvironment creates a distributional gap that risks hallucinating precision.

These practical obstacles are compounded by a more fundamental conceptual gap. A true digital twin, as defined by the National Academies of Sciences, Engineering, and Medicine (NASEM), requires dynamic, bidirectional interaction between the digital model and the physical system it represents. Clinical oncology rarely affords this: patient data typically arrives as sparse, static snapshots—a biopsy at diagnosis, perhaps another at progression—rather than the iterative feedback loop the definition demands. Even a technically perfect *ab initio* model would therefore struggle to qualify as a genuine digital twin without the infrastructure for iterative updating. This double bind—models that are both hard to build and hard to keep current—motivates a fundamentally different strategy.

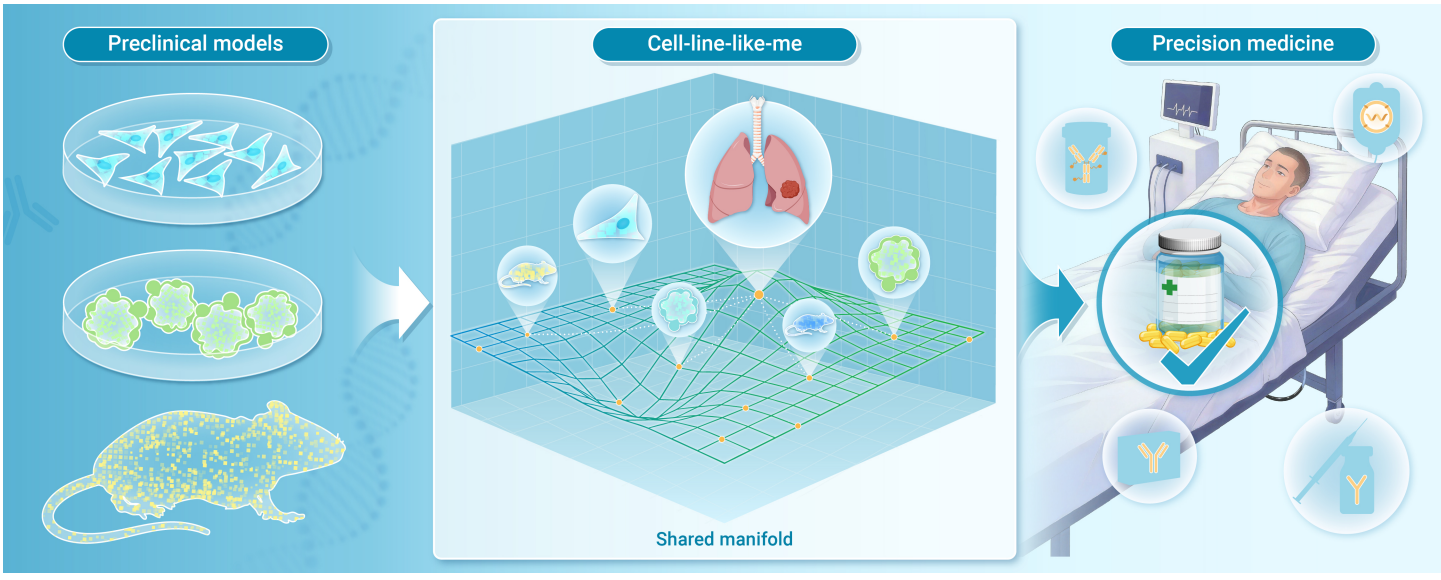

Figure 1. Conceptual overview of the Cell-Line-Like-Me framework (1) Patient tumors and preclinical models (cell lines, PDX, organoids) are independently profiled across multi-omics modalities. (2) A shared biological manifold is constructed via manifold learning, separating shared biological coordinates from domain-specific distortions. (3) Drug response is inferred geometrically from the patient's coordinate neighborhood among drug-response-annotated preclinical anchors, enabling patient-specific therapeutic prioritization in a precision medicine setting.

THE MANIFOLD HYPOTHESIS: A GEOMETRY OF BIOLOGY

If we accept that cancer models are imperfect yet contain "islands of truth", and we acknowledge the risks of *ab initio* simulation, a practical third path emerges: we do not need to reconstruct the cell from nothing and can instead learn the geometry of biological state space and work within it. This is the core of the manifold hypothesis.

The statistical intuition is not new. Dimensionality reduction has long shown that high-dimensional expression data admits low-rank representations; the modern manifold framing simply relaxes the linearity assumption, allowing biological states to lie on curved surfaces recoverable from local neighborhood structure. Biology has carried a parallel intuition: Waddington's epigenetic landscape and later dynamical-systems views treat cell phenotypes as attractors within a structured state space. Whether cast as latent factors or manifolds, the claim is consistent: the intrinsic dimension of biologically valid states is far smaller than the dimension we measure.

This insight motivates us to propose the concept named "Cell-Line-Like-Me": learning a shared biological manifold from both preclinical models and patient tumors, then embedding a patient's tumor as a coordinate on this manifold. Because preclinical models occupying nearby positions carry dense drug-response annotations, the patient's therapeutic sensitivity can be inferred geometrically from its neighborhood, transforming translational prediction from gene-level matching into a coordinate-based inference problem. Computationally, the shared manifold is constructed by learning a low-dimensional embedding in which domain-specific distortions are penalized while biologically relevant variation is preserved.

"CELL-LINE-LIKE-ME": MAPPING COORDINATES, NOT JUST MATCHING GENES

In this framework, a patient's digital representation is a coordinate on a shared manifold co-defined by patient tumors and preclinical models. The core advantage comes from "labeled anchors": large cancer model resources carry dense drug-response profiles, whereas patients are effectively unlabeled at decision time. Once both are embedded in a comparable space, drug response prediction becomes a geometric problem: inferring sensitivity from the patient's neighborhood among labeled anchors.

However, embedding tumors and models into a shared space is necessary but not sufficient. If the alignment is only global, matching overall expression profiles without resolving which biological axes are truly shared, the framework inherits the same fallacy that has plagued preclinical translation for decades. The clinical record makes this danger concrete. A systematic analysis of insulin-like growth factor 1 receptor (IGF-1R) inhibitors documented 183 clinical trials and billions in aggregate R&D costs, yet near-universal failure driven in part by resistance mechanisms present *in vivo* but absent from the cell-line models that motivated development.[5] The models appeared globally similar to the patient tumors they were meant to represent yet diverged precisely along the mechanistic axes that determined drug response. The lesson is clear: manifold alignment must be pathway-aware, ensuring that proximity in the shared space reflects congruence along the biological dimensions that matter for a given drug, rather than a misleading aggregate resemblance.

A critical design choice is which preclinical systems serve as anchors—a question with both a vertical and a horizontal dimension. Vertically, candidate therapies traverse a hierarchy of experimental systems before reaching patients: *in vitro* models (cell lines, organoids) provide the densest drug-response annotations but lack stroma, vasculature, and immune context; murine *in vivo* models (CDX, PDX, GEMM) restore three-dimensional architecture and pharmacokinetics but introduce cross-species divergences. Each tier carries a different congruence profile with human tumors. The framework is tier-agnostic: the practical starting point is the cancer cell line, because CCLE, GDSC, and CTRP collectively provide matched genomic, transcriptomic, and pharmacological profiles for over a thousand lines, but PDX and organoid databases can be progressively incorporated as additional anchor layers. Horizontally, different biological axes are preferentially revealed by different data modalities: bulk RNA-seq is the most available default, but whole-exome sequencing captures mutational drivers, Cytometry by Time-Of-Flight (CyTOF) quantifies immune composition invisible to transcriptomics, and phosphoproteomics measures signaling activation at the post-translational level. The framework should therefore be conceived as a family of modality-specific or multi-modal manifolds suited to different therapeutic questions.

Building the shared space can follow complementary routes: direct harmonization of batch and platform effects (e.g., ComBat, Celligner) ; linear factor models yielding interpretable shared coordinates (e.g., NMF, LIGER); and nonlinear manifold learning via deep generative models or adversarial domain adaptation (e.g. scVI). Critically, the framework must also reject poor surrogates rather than always forcing a match--through manifold distance thresholding, pathway-specific congruence scores, and benchmarking against held-out clinical outcomes.

FROM REPLICAS TO COORDINATES

This commentary reframes what a tumor "digital twin" should prioritize today. The *ab initio* ideal, reconstructing a patient tumor as a fully specified simulator, remains aspirational in the face of identifiability limits, technical confounding, and domain shift. A more immediately testable goal is to redefine the twin as a coordinate system problem: learn a shared biological manifold in which tumors and models become comparable along the dimensions that matter, then connect patients to labeled anchors with explicit control of nuisance variation and clear validation criteria.

Cell-Line-Like-Me seeks to provide a unifying conceptual bridge linking manifold-based alignment, pathway-aware congruence, and multi-tier model selection into a single translational framework grounded in digital twin principles. Validating this framework requires a three-step workflow: (1) internal cross-validation of drug response prediction within the preclinical anchor space; (2) external benchmarking against held-out clinical outcomes from independent patient cohorts; and (3) pathway-specific congruence auditing to verify that the biological axes driving alignment are mechanistically relevant to the therapeutic target.

Important limitations remain. When anchor models are sparse, as for rare cancer types, manifold estimates become unreliable, and the framework must explicitly flag low-confidence predictions via distance thresholds rather than forcing a match. Clinical deployment further requires prospective validation, regulatory clarity for AI-assisted treatment selection, and integration with existing decision-support workflows. Finally, the entire approach depends on robust cross-platform harmonization, an area of active methodological development.

By converting decades of preclinical drug-response knowledge into patient-specific hypotheses that are geometrically interpretable, uncertainty-aware, and experimentally falsifiable, it offers a scalable path toward a digital twin that earns its name—not by claiming perfect replication, but by delivering reliable, mechanism-linked guidance where congruence is real and measurable.

REFERENCES

1. Zong W., Rahman T., Zhu L., et al. (2023). Transcriptomic congruence analysis for evaluating model organisms. *Proc. Natl. Acad. Sci. USA* **120**:e2202584120. DOI:10.1073/pnas.2202584120
2. Mourragui S., Loog M., van de Wiel M.A., et al. (2019). PRECISE: A domain adaptation approach to transfer predictors of drug response from pre-clinical models to tumors. *Bioinformatics* **35**:i510–i519. DOI:10.1093/bioinformatics/btz372
3. Zou J., Shah O., Chiu Y.C., et al. (2024). Systems approach for congruence and selection of cancer models towards precision medicine. *PLoS Comput. Biol.* **20**:e1011754. DOI:10.1371/journal.pcbi.1011754
4. Laubenbacher R., Mehrad B., Shmulevich I., et al. (2024). Digital twins in medicine. *Nat. Comput. Sci.* **4**:184–191. DOI:10.1038/s43588-024-00607-6
5. Jentzsch V., Osipenko L., Scannell J.W., et al. (2023). Costs and causes of oncology drug attrition with the example of insulin-like growth factor-1 receptor inhibitors. *JAMA Netw. Open* **6**:e2324977. DOI:10.1001/jamanetworkopen.2023.24977

FUNDING AND ACKNOWLEDGMENTS

This work was supported by the Chongqing Municipal Science and Technology Bureau (Grant No. CSTB2024YCJH-KYXM0098) and the Venture and Innovation Support Program for Chongqing Overseas Returnees (Grant No. 203011220240002). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare no competing interests.