OPEN ACCESS

Digital Medicine PERSPECTIVE

The Innovation Informatics

Digital twins in medicine: Future horizons

Mengbi Shen,[1] Shiyu Du,[2] and Jian He[1,*]

[1]Center for Single-Cell Omics, School of Public Health, Shanghai Jiao Tong University School of Medicine, Shanghai 20027, China

[2]Qingdao Institute of Software, College of Computer Science and Technology, China University of Petroleum (East China), Shandong Key Laboratory of Intelligent Oil & Gas Industrial Software, Qingdao 266580, China

*Correspondence: jih003@sjtu.edu.cn

Received: February 11, 2026; Accepted: August 18, 2026; Published Online: August 18, 2026; https://doi.org/10.59717/j.tii.2027.100078

© 2027 The Author(s). This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

Citation: Shen M., Du S. and He J. (2027). Digital twins in medicine: Future horizons. *The Innovation Informatics* **3**:100078.

Virtual patients have advanced from proof-of-concept to pragmatic tools in precision medicine, yet their role must be recalibrated. This Perspective argues that virtual patients should function not as substitutes for human participants, but as ethically governed auxiliary systems for addressing data scarcity, high-dimensional multi-omics complexity, and clinical trial constraints. Their value lies less in statistical mimicry than in mechanistically interpretable links between interventions and physiological responses. This requires hybrid architectures that integrate biological knowledge, pathway logic, PK/PD principles, and adaptive learning. A major bottleneck remains validation: retrospective consistency is insufficient, and prospective testing against independent real-world cohorts with pre-specified endpoints and equivalence margins is needed. Virtual patients' highest near-term utility resides in pre-trial simulation—optimizing dosing, refining stratification, and mitigating late-stage failure risks—not in supplanting phase III RCTs. Ultimately, virtual patients should work symbiotically with empirical studies, exploring counterfactuals in silico while reserving definitive validation for real-world evidence.

CURRENT LANDSCAPE AND POSITIONING

Virtual patients are steadily transcending their initial proof-of-concept phase and entering the arena of practical clinical research. Yet their significance should not be romanticized as mere digital replicas of flesh-and-blood individuals; rather, they ought to be thoughtfully repositioned as a pragmatic supplement—one emerging precisely at the intersection of scarce biospecimens, protracted trial timelines, and stringent ethical constraints. In precision medicine, researchers confront a paradox: high-dimensional multi-omics data are expensive and yield few samples. Yet analytical ambitions demand modeling within feature spaces that span tens of thousands of molecular dimensions.[1] This "small-n, high-dimensionality" conundrum renders conventional statistical inference perilously vulnerable to overfitting, while simultaneously stalling promising biomarker discovery at the very threshold of insufficient statistical power. For example, in rare molecularly defined cancers, virtual cohorts may help explore plausible treatment-response patterns and refine enrichment criteria before these hypotheses are tested in external patient cohorts.

In response to these structural limitations, synthetic data generation and virtual cohort expansion have gradually evolved from theoretical curiosities into strategic necessities. Still, a critical caveat must be emphasized: virtual patients should never be misconstrued as wholesale substitutes for human participants in clinical investigation. A more grounded and ethically sound trajectory positions them as an augmentation layer—complementing rather than replacing real-world trials. Practical use cases include synthetic control arms to reduce recruitment pressure, virtual simulations to examine heterogeneous responses across molecular subgroups, and computational screening to identify potential hyper-responders or patients at elevated risk of adverse events before trial enrollment. Generative approaches leveraging GANs or VAEs can indeed broaden sample distributions at a statistical level, but whether their outputs encode biologically interpretable mechanisms or support robust causal inference remains an open and pressing question.[2] Herein lies a crucial epistemological boundary often glossed over in waves of technological optimism: statistical similarity does not equate to mechanistic equivalence. Two datasets may appear indistinguishable in distributional metrics while embodying fundamentally divergent biological realities.

METHODOLOGICAL FRAMEWORKS AND VALIDATION

The foundational challenge in constructing trustworthy virtual patients, therefore, resides in navigating a delicate and dynamic equilibrium between simplification and fidelity. Multimodal clinical data—from genomic sequences and proteomic profiles to radiological images and electronic health records—must be distilled into compact yet evolvable state representations. Yet an obsessive pursuit of granular realism frequently triggers model complexity explosions that paradoxically undermine both interpretability and cross-cohort generalizability. The true litmus test lies not in how precisely a model reconstructs historical observations, but in its capacity to encode causal pathways linking interventions to physiological responses. A mature virtual patient model should reliably distinguish state transitions driven by natural disease progression from those induced by therapeutic perturbations—whether a kinase inhibitor halting tumor growth or an immunomodulator reshaping the tumor microenvironment. This represents a decisive departure from conventional machine learning's correlational paradigm, pivoting instead toward mechanistic modeling of disease dynamics as an evolving biological process. Artificial intelligence's highest value here extends beyond pattern recognition; it should serve as a bridge integrating biophysical principles, pathophysiological knowledge, and data-driven learning to construct in silico experimental platforms capable of simulating intervention consequences under controlled conditions. Beyond oncology, similar challenges arise in domains with limited sample sizes. For instance, cardiovascular digital twins integrating hemodynamic models show potential in predicting trajectories.[3] These examples underscore that the value of virtual patients extends beyond oncology, particularly in contexts where "small-n" data fundamentally limits conventional statistical inference. This integrative philosophy provides essential methodological grounding for building virtual patients that are not merely statistically plausible but biologically meaningful.[4]

Nevertheless, the absence of a rigorous validation loop remains the field's most significant bottleneck. The majority of current studies still anchor their credibility in internal consistency checks or retrospective validation against historical datasets—approaches inherently vulnerable to confirmation bias and temporal confounding. What is urgently needed are prospective, head-to-head comparisons against independent real-world cohorts, with pre-specified equivalence margins and transparent reporting standards. In practice, this means that the model, input variables, endpoints, and analysis plan should be locked before prospective testing. Predictions should then be compared with observed outcomes in an independent cohort before any post hoc model updating is performed. A principle underscored by recent oncology AI translation efforts where even FDA-cleared tools require prospective randomized trials to demonstrate impact on hard endpoints rather than relying on retrospective metrics alone.[5] Regulatory agencies have begun laying groundwork through methodological qualification frameworks, yet the research community must maintain sober recognition: virtual patients are, by design, inference engines built upon specific mechanistic assumptions and data distributions.[6] Their highest utility does not reside in replacing phase III randomized controlled trials—the gold standard for therapeutic validation—but in functioning as simulation sandboxes during the trial design phase. Before enrolling a single human participant, researchers could leverage virtual cohorts to explore alternative dosing regimens, evaluate competing clinical endpoints, or refine patient stratification criteria. Such pre-trial simulation holds genuine promise for optimizing statistical power, reducing sample size requirements, and mitigating the staggering financial and human costs associated with late-stage trial failures. This pragmatic, risk-mitigation role may

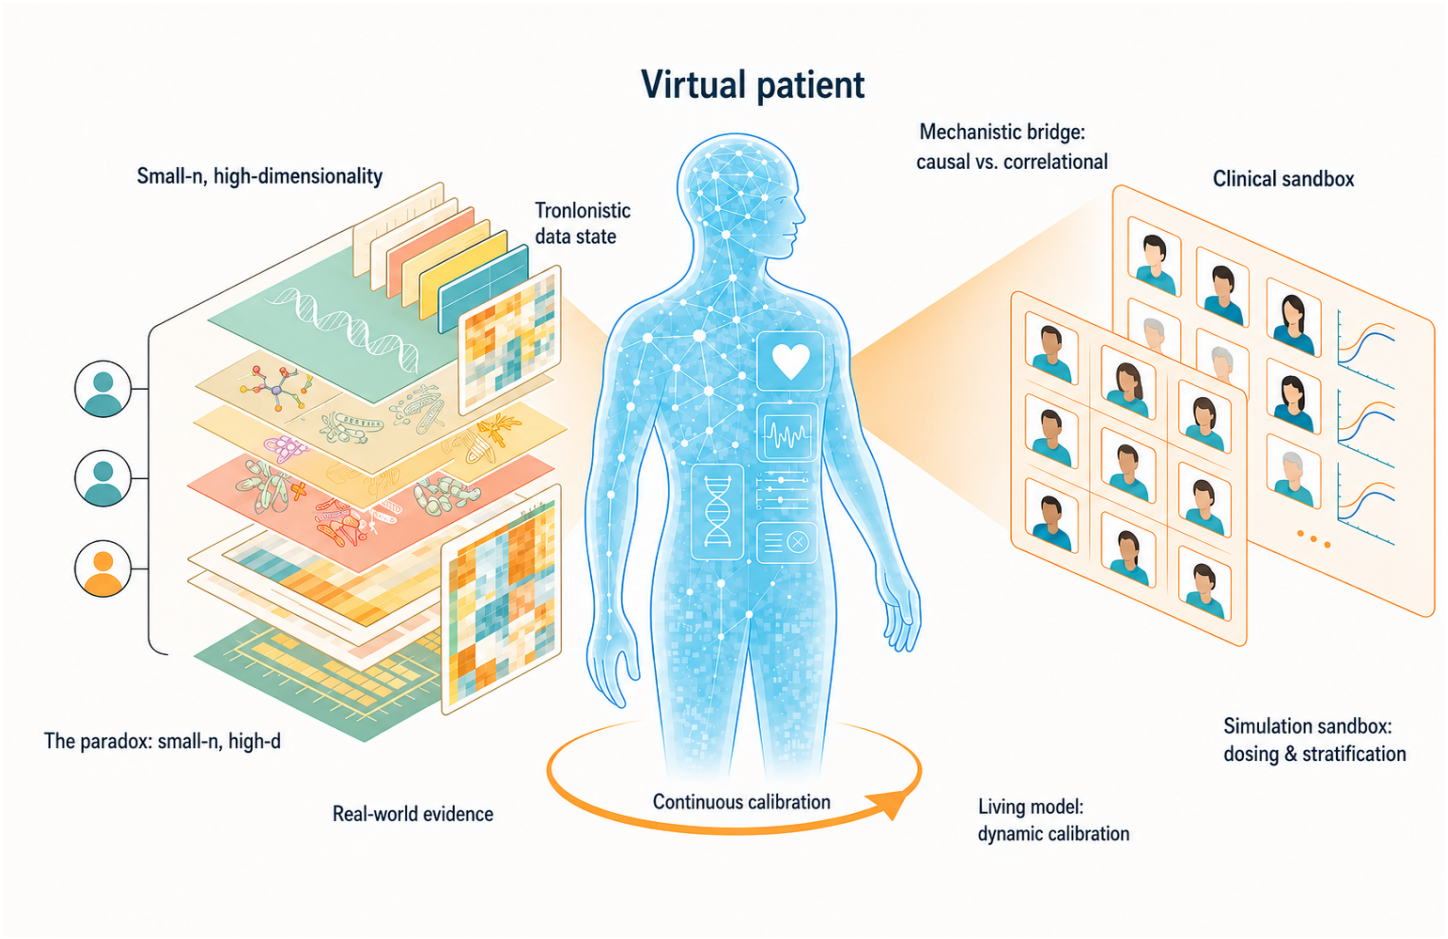

Figure 1. Virtual patients as dynamically calibrated digital twins for precision medicine Small-sample, high-dimensional multi-omics and clinical data are integrated to construct virtual patients that link patient-state representations with mechanistic disease processes. Through continuous calibration with real-world evidence, these models can support clinical and simulation sandboxes for optimizing dosing, refining patient stratification, and improving trial design.

well represent virtual patients' most tangible and ethically defensible near-term contribution to clinical research.

FUTURE PERSPECTIVES

Maturation hinges less on algorithmic novelty than on forming closed-loop systems with living entities. Unlike static predictive models deployed once and frozen in time, virtual patients more closely resemble "online evolving models" whose internal representations must update iteratively as new real-world evidence accumulates—longitudinal follow-up data, dynamic biomarker trajectories captured through liquid biopsies, or granular treatment response patterns emerging from real-world evidence platforms. In addition to clinical trial data, large-scale natural population studies provide a critical resource for establishing baseline physiological variability. Integrating these cohorts into virtual patient frameworks could extend their utility beyond interventional simulation to support early disease warning, helping to identify at-risk trajectories before clinical manifestation. Only within such a perpetual "data–model–revalidation" cycle can virtual patients gradually approximate authentic disease trajectories rather than remaining frozen as static statistical artifacts. Conceptually, we must shift from viewing virtual patients as discrete objects to be engineered and delivered, toward recognizing them as process-oriented systems requiring continuous refinement through iterative dialogue with clinical reality.

Methodologically, the field is witnessing a decisive pivot from purely data-driven paradigms toward hybrid architectures that thoughtfully blend "mechanistic constraints + adaptive learning." Deep learning excels at uncovering patterns in high-dimensional spaces, integrating multi-omics and imaging.[7] Yet without anchoring in established biological or pathophysiological principles, such models often yield inscrutable "black box" predictions that resist clinical interpretation and fail to generalize beyond training distributions. Responding to this limitation, an emerging wave of research now strategically embeds domain knowledge—metabolic network topologies, signaling pathway logic, cellular differentiation trajectories, or pharmacokinetic–pharmacodynamic (PK/PD) formalisms—directly into virtual patient architectures. The aspiration transcends superficial data mimicry; the goal is to construct models that not merely "resemble data" in statistical distribution, but "resemble disease" in causal structure and dynamic behavior. This transition from correlation to causation marks a watershed moment in elevating virtual patients from engineering conveniences to genuine scientific instruments.

Notably, the conceptual scale of virtual patients is undergoing profound re-examination. Early implementations focused predominantly on population-level risk prediction—averaging heterogeneity into smoothed epidemiological curves. Contemporary advances, however, trend toward multi-scale integration spanning "cellular–tissue–organ–individual" hierarchies. Consider a virtual patient resting on single-cell states from biopsies, mapped onto tissue architectures, then scaled to predict organ dysfunction and outcomes. This bottom-up digital twin pathway transforms virtual patients from statistical abstractions into biologically grounded digital entities that preserve molecular heterogeneity and retain capacity for dynamic evolution. Such granularity offers structurally richer explanatory frameworks for phenomena that confound conventional risk models—why certain patients develop rapid therapeutic resistance, why immune-related adverse events manifest unpredictably, or how microenvironmental niches foster metastatic seeding.

Yet as virtual patients approach greater mechanistic fidelity, their inherent engineering boundaries become increasingly apparent. Model complexity inevitably inflates parameter spaces at a rate far outpacing the accumulation of high-quality calibration data. This asymmetry unavoidably imbues virtual patients with a "hypothesis-laden" character: every architectural choice—from which pathways to include to how cellular interactions are parameterized—implicitly encodes prior assumptions about disease biology.[8] Acknowledging

this reality, the field would benefit from shifting emphasis away from the elusive goal of "perfect digital replication" toward cultivating "problem-oriented minimal viable models." This "fit-for-purpose" philosophy deliberately scopes models to address specific clinical questions—such as identifying hyper-responders to checkpoint inhibitors or optimizing dosing in renally impaired subpopulations—while resisting the seduction of ontological completeness.[9] A practical applicability framework should therefore specify three elements: the clinical decision being supported, the minimum biological resolution required, and the evidence threshold needed before model outputs influence trial design or clinical interpretation. Beyond methodological considerations, ethical governance should be operationalized through transparent documentation of data provenance, consent scope, privacy protection, subgroup performance, bias assessment, and model-update rules, thereby helping safeguard patient safety, equity, and reproducibility.

CONCLUSION

In summation, virtual patients are not replacements for biological organisms, but rational complements in settings marked by data scarcity, mechanistic uncertainty, and ethical constraints. Their enduring value lies in clarifying intervention pathways, supporting iterative calibration with real-world evidence, and improving trial design before patients are exposed to unnecessary risk.[10] To realize this value, the field must move beyond retrospective plausibility toward prospective validation, transparent reporting, and fit-for-purpose applicability standards. The mature role of virtual patients will therefore be symbiotic: simulation can explore counterfactuals and refine hypotheses, while empirical studies remain essential for definitive evidence of safety and clinical benefit.

REFERENCES

1. Iqbal J.D., Krauthammer M., Witt C.M., et al. (2025). A consensus statement on the use of digital twins in medicine. *Npj Digit. Med.* **8**:484. DOI:10.1038/s41746-025-01897-4
2. He W., Wu X., Jin Z., et al. (2025). Generative artificial intelligence in medical imaging: Current landscape, challenges, and future directions. *Interdiscip. Med.* **3**:e20250024. DOI:10.1002/INMD.20250024
3. Geddes J.R., Jensen C.W., Tanade C., et al. (2025). Digital twins for noninvasively measuring predictive markers of right heart failure. *Npj Digit. Med.* **8**:545. DOI:10.1038/s41746-025-01920-8
4. Xu Y., Liu X., Cao X., et al. (2021). Artificial intelligence: A powerful paradigm for scientific research. *The Innovation* **2**:100179. DOI:10.1016/j.xinn.2021.100179
5. Lotter W., Hassett M.J., Schultz N., et al. (2024). Artificial intelligence in oncology: Current landscape, challenges, and future directions. *Cancer Discov.* **14**:711–726. DOI:10.1158/2159-8290.CD-23-1199
6. Bai L. and Su J. (2026). Artificial intelligence virtual organoids (AIVOs). *Bioact. Mater.* **59**:45–68. DOI:10.1016/j.bioactmat.2025.12.030
7. Katsoulakis E., Wang Q., Wu H., et al. (2024). Digital twins for health: A scoping review. *Npj Digit. Med.* **7**:77. DOI:10.1038/s41746-024-01073-0
8. De Domenico M., Allegri L., Caldarelli G., et al. (2025). Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *Npj Digit. Med.* **8**:37. DOI:10.1038/s41746-024-01402-3
9. Bunne C., Roohani Y., Rosen Y., et al. (2024). How to build the virtual cell with artificial intelligence: Priorities and opportunities. *Cell* **187**:7045–7063. DOI:10.1016/j.cell.2024.11.015
10. Wang H., Arulraj T., Ippolito A., et al. (2024). From virtual patients to digital twins in immuno-oncology: Lessons learned from mechanistic quantitative systems pharmacology modeling. *Npj Digit. Med.* **7**:189. DOI:10.1038/s41746-024-01188-4

FUNDING AND ACKNOWLEDGMENTS

This work was supported by the National Key Research and Development Program of China (2022YFD2101503), Advanced Materials-National Science and Technology Major Project (2026ZD0623600), the National Key R&D Program of China (2024YFB3817300) and Machine learning-based drug repositioning for tumor-related targets (2024HX012). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

M. Shen contributed to the methodology and writing - original draft. J. He contributed to the conceptualization and supervision. J. He and S. Du contributed to riting - review & editing and funding acquirment. All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare no competing interests.

ETHICAL STATEMENT AND PATIENT CONSENT

Not applicable.