

Catalyzing CAR-T innovation with AI virtual cells: Breakthroughs, bottlenecks, and the path forward

Xiaojie Liang,¹ Baiwei Luo,² Ning Jiang,^{3,4,5} and Hua-Bing Li^{6,7,8,*}

¹Department of Hematology, The First Affiliated Hospital of Chongqing Medical University, Chongqing 400016, China

²Department of Hematology, Nanfang Hospital, Southern Medical University, Guangzhou 510515, China

³Department of Pathology, College of Basic Medical Sciences, Chongqing Medical University, Chongqing 400016, China

⁴Molecular Medicine Diagnostic and Testing Center, Chongqing Medical University, Chongqing 400016, China

⁵Department of Pathology, The First Affiliated Hospital of Chongqing Medical University, Chongqing 400016, China

⁶Institute for Immunology and Pathogenesis (IIP), Chongqing Medical University, Chongqing 400016, China

⁷Center for Immune-Related Diseases at Shanghai Institute of Immunology, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai 200025, China

⁸Department of Immunobiology, Yale University School of Medicine, New Haven CT 06520-8055, USA

*Correspondence: huabing.li@shsmu.edu.cn

Received: February 26, 2026; Accepted: September 29, 2026; Published Online: September 30, 2026; <https://doi.org/10.59717/j.xinn-life.2026.100261>

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

Citation: Liang X., Luo B., Jiang N., et al. (2027). Catalyzing CAR-T innovation with AI virtual cells: Breakthroughs, bottlenecks, and the path forward. *The Innovation Life* 5:100261.

The emergence of artificial intelligence virtual cell (AIVC) technology represents a new paradigm in biomedical research.¹ Chimeric antigen receptor T-cell (CAR-T) therapy is established for selected hematologic malignancies and is being explored in autoimmune diseases.² On April 10, 2025, the U.S. Food and Drug Administration (FDA) announced a roadmap to reduce and potentially replace selected animal-testing requirements using validated new approach methodologies, including AI-based computational models, cell-based assays, and organoids. Against this backdrop, integrating AIVCs into CAR-T research may accelerate development. This Commentary highlights emerging AIVC applications in CAR-T therapy (Figure 1) and persistent technical challenges.

AIVCs: CONCEPTUAL FRAMEWORK AND CURRENT LANDSCAPE

AIVCs are AI-enabled computational frameworks that integrate heterogeneous biological information across relevant scales and modalities with multiple computational paradigms, including large-scale neural networks, mechanistic and mathematical models, agent-based models, graph-based approaches, and other data-driven methods, to represent, simulate, and predict cellular states and responses.¹ Inputs include DNA, RNA, and protein sequences; single-cell multi-omics; imaging; spatial omics; and perturbation data, encoded by architectures such as transformers, autoencoders, vision models, and graph neural networks. These data are mapped into interconnected universal representations (URs) of molecules, cells, and tissues. Virtual instruments (VIs) operate on URs: decoder VIs translate latent representations into interpretable phenotypes or molecular readouts, whereas manipulator VIs model transitions between biological states.

A recent whole-cell digital-twin framework integrates live-cell 4D lattice light-sheet microscopy with particle-based reaction-diffusion simulations to reproduce mitochondrial and microtubule dynamics and predict responses to cytoskeletal perturbations.³ CAR-T-specific AI modeling, exemplified by CAROT, addresses CAR design and single-cell responses, but an integrated CAR-T AIVC remains to be established and rigorously validated.⁴

AIVCs remain emerging rather than fully realized digital-cell frameworks. Their potential advantage over conventional PK/PD and QSP models, ODE- and agent-based models, and machine-learning approaches is the integration of multimodal cellular representations, heterogeneous cell states, and mechanistic constraints. A recent benchmark of single-cell genetic perturbation prediction found that several deep-learning models did not outperform simple linear baselines.⁵ AIVCs should thus be considered complementary and potentially hybrid frameworks, with CAR-T applications remaining translational hypotheses requiring rigorous validation.

POTENTIAL APPLICATIONS OF AIVCS IN CAR-T DEVELOPMENT

Target discovery

Resistance and relapse remain major challenges in CAR-T therapy, creating a need for next-generation products. However, conventional CAR-T development is protracted and resource-intensive.

Predicting cellular behavior is a key feature of AIVCs. By learning from large-scale datasets spanning spatially and temporally resolved, perturbed, and unperturbed conditions, AIVCs may capture CAR-T-associated dynamics across molecular, cellular, and tissue levels and model interactions between CAR-T and tumor cells. AIVCs could virtually screen candidate genes by

predicting effects on tumor-cell killing, thereby prioritizing targets for experimental validation. More mature frameworks could compare single-, dual-, or multi-target CAR-T designs while modeling baseline antigen heterogeneity and treatment-induced antigen-state changes. By estimating persistence or outgrowth of target-low or target-negative states after virtual perturbation, they may prioritize combinations with broader tumor-cell coverage and lower vulnerability to antigen escape. Whether AIVC-guided target selection reduces antigen-escape-associated relapse or improves durable clinical outcomes requires prospective validation.

Resistance mechanism prediction

Treatment resistance in CAR-T therapy is driven by tumor heterogeneity and evolution; T-cell exhaustion and an immunosuppressive microenvironment also contribute to treatment failure.

The added value of AIVCs is not to rediscover these mechanisms but to integrate them into dynamic, context-specific response models. By combining generative or state-transition models constrained by mechanistic knowledge, future AIVC frameworks could generate virtual tumor populations with diverse molecular phenotypes and functional states. These frameworks could integrate longitudinal and perturbational data to simulate state transitions and reciprocal interactions among malignant cells, CAR-T cells, and accessory immune populations under therapeutic pressure. These models may prioritize testable interventions, such as alternative or multi-target CAR designs. However, AIVC-derived predictions should remain hypothesis-generating until independently and prospectively validated.

Toxicity forecasting and mitigation

A major challenge in CAR-T therapy is the risk of severe toxicities, such as cytokine release syndrome (CRS) and immune effector cell-associated neurotoxicity syndrome (ICANS).

AIVCs could, in principle, help address this gap by enabling quantitative simulation of immune cascades and in silico screening of intervention strategies. Specifically, AIVCs could be trained on large-scale longitudinal and time-resolved datasets to infer CAR-T cell-state trajectories and dynamic interaction networks, which may subsequently support the prediction of cytokine-release patterns over time. In agent-based modeling (ABM), each virtual cell follows predefined, often probabilistic biological rules that map intrinsic states and local signals to behaviors; system-level dynamics then emerge from intercellular interactions. Activation thresholds, transition probabilities, and secretion/decay rates should be experimentally informed or independently calibrated, and sensitivity analyses should assess robustness. After rigorous validation, such models could support quantitative investigation of selected CRS/ICANS dynamics.

Researchers could also test intervention timing, dosage, and drug combinations for reducing simulated cytokine surges while assessing effects on CAR-T antitumor efficacy. Personalization would require calibration with patient-specific disease and antigen features, tumor and immune-cell molecular profiles, CAR-T product characteristics, baseline clinical variables, and longitudinal CAR-T, cytokine, toxicity, and response data. Clinical translation will require multicenter external validation and prospective evaluation across diseases, populations, and CAR-T products.




Figure 1. Conceptual framework for AI virtual cells to catalyze CAR-T innovation Multimodal multi-omic, spatial and longitudinal, perturbational, and CAR-T/clinical data are harmonized into multiscale representations. AI architectures are integrated with biological priors to construct AI virtual cells (AIVCs), which are calibrated to patient-specific or experimental contexts and used to perform virtual simulations and predict cell-state transitions, resistance trajectories, CAR-T fitness, cytokine dynamics, treatment phenotypes, and associated uncertainty. The lower panel summarizes four potential applications of AIVCs in CAR-T development: target discovery, resistance mechanism prediction, toxicity forecasting and mitigation, and precision treatment. Model predictions are iteratively evaluated through experimental and clinical validation, forming a feedback loop for continued refinement.

Precision treatment

CAR-T responses vary substantially among patients, and treatment-related toxicities remain common, underscoring the need for integrated multidimensional biological profiling.

The preceding applications address individual predictive tasks, whereas integrative patient-specific decision support represents a higher-level synthesis of their outputs. A patient-calibrated AIVC could eventually compare candidate strategies under predefined efficacy, resistance, and toxicity criteria. With sufficient prospective validation, such a model could serve as a patient-specific computational surrogate. However, no AIVC has yet been prospectively validated for individualized CAR-T decision-making.

CURRENT CHALLENGES

The construction of AIVCs involves several major challenges: 1) High-dimensional and large-scale data. Training requires representative datasets across temporal and spatial scales. For CAR-T research, these should include perturbation-resolved, longitudinal, and spatial multi-omics of CAR-T cells, malignant cells, and accessory immune cells with explicit clinical annotation. Such datasets remain scarce because serial sampling is constrained by specimen availability, cost, patient condition, and tissue accessibility. Multi-center consortia should therefore harmonize sampling, processing, clinical endpoints, CAR-T product metadata, and interoperable formats. 2) Interpretability and validation. The black-box nature of AI models remains a barrier to clinical application. Interpretability should be designed into CAR-T AIVCs. Counterfactual perturbation and explanation-stability analyses could help determine whether predictions rely on biologically plausible and reproducible signals. Additionally, no AIVC should be assumed to reproduce all relevant biology. Validation should proceed hierarchically: comparison with simple baselines on held-out and out-of-distribution datasets; external validation across patients, centers, and assay platforms; prospective testing in co-culture, organoid, or other human-relevant experimental systems; and, where necessary, in vivo and clinical validation. Calibration, uncertainty, biological fidelity, and decision-relevant performance should be reported before AIVCs are considered for clinical decision support. 3) Ethics and safety. Personalized AIVCs use sensitive genomic, molecular, product, and longitudinal clinical data. Governance should include informed consent, purpose-limited access, security safeguards, provenance and audit trails, and oversight of secondary use and model updating. 4) Resource requirements. Longitudinal multi-omics generation, secure data harmonization, high-performance computing, repeated simulation, and model maintenance may be inaccessible to many institutions. Open pretrained models, shared infrastructure, parameter-efficient adaptation, and model compression may improve accessibility, but computational cost should be reported alongside predictive performance and

justified by demonstrable incremental value.

Several additional barriers limit translation of AIVCs to CAR-T development. Inadequate anatomical fidelity constrains their ability to replace animal models or early-phase clinical trials. AIVCs may oversimplify cellular heterogeneity and spatiotemporal immune complexity, requiring rigorous in vivo validation of predicted targets and interventions. Generalizability across populations and disease contexts

may also be limited, creating risks of algorithmic bias in underrepresented groups and necessitating subgroup-specific evaluation and external validation.

A pragmatic roadmap should be staged. In the short term, the field needs harmonized CAR-T-specific longitudinal datasets and transparent benchmarks against simpler models. The medium-term priority is to develop interpretable hybrid models and validate their predictions in orthogonal human-relevant systems. Long-term translation will require interdisciplinary collaboration, prospective evaluation within predefined contexts of use, uncertainty reporting, regulatory oversight, and evidence of incremental value. AIVCs should ultimately be judged by reproducible utility rather than model scale or novelty alone.

REFERENCES

- 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
- Rafei H, Upadhyay R, and Sharma P. (2026). A guide to CAR T cell therapies: Development, current status and future prospects. *Nat. Rev. Immunol.* DOI:10.1038/s41577-026-01322-1
- Arkfeld E, Wang Z, Hakozi H, et al. (2026). Whole-cell particle-based digital twin simulations from 4D lattice light-sheet microscopy data. *Cell* **189**:4581–4593.e6. DOI:10.1016/j.cell.2026.06.010
- Driessen A, Born J, Castellanos Rueda R, et al. (2026). Modeling chimeric antigen receptor response at the single-cell level with conditional optimal transport. *Cell Syst.* **17**:101591. DOI:10.1016/j.cels.2026.101591
- Ahlmann-Eltze C, Huber W, and Anders S. (2025). Deep-learning-based gene perturbation effect prediction does not yet outperform simple linear baselines. *Nat. Methods* **22**:1657–1661. DOI:10.1038/s41592-025-02772-6

FUNDING AND ACKNOWLEDGMENTS

This work was supported by the Ministry of Science and Technology of China (2021YFA1100800), the National Natural Science Foundation of China (82325024, 82341017, 82350112, 82300342, 82441048 and 82461160323). Figure 1 was created by BioRender (<https://www.biorender.com/>). 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.