



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Pharmaceutical Biotechnology Needs Digital Twins That Preserve Biological Uncertainty

Thabo Nkosi^{1*}, Lerato Molefe¹, Siphon Dlamini², Ayanda Mokoena³, Kabelo Ndlovu⁴

¹Department of Pharmaceutical Digital Twins and Biological Uncertainty, Faculty of Pharmacy, University of Cape Town, Cape Town, South Africa.

²Department of Bioprocess Modeling and Simulation, Faculty of Pharmacy, University of the Witwatersrand, Johannesburg, South Africa.

³Department of Uncertainty Preservation in Digital Twins, Faculty of Pharmaceutical Sciences, University of Pretoria, Pretoria, South Africa.

⁴Department of Biotechnology Informatics, Faculty of Pharmacy, Stellenbosch University, Stellenbosch, South Africa.

***Email:** thabo.nkosi@uct.ac.za

ABSTRACT

Digital twins are increasingly positioned as enabling technologies for model-based pharmaceutical biotechnology, yet many proposed architectures treat biological uncertainty as a nuisance to be averaged away rather than as part of the state that must be represented. This distinction matters because a bioprocess may remain numerically predictable while its underlying cell population, model structure, sensor validity, or scale-dependent behavior changes in ways that alter the meaning of that prediction. This Methodological Perspective argues that a useful bioprocess digital twin should therefore represent not only estimated process states but also uncertainty about biological state, parameters, model structure, data provenance, and model validity. Mechanistic and data-driven models are treated as complementary rather than interchangeable components; soft sensors are separated from direct observations; and state updating is distinguished from evidence that the governing model remains adequate. The article develops a proposed probabilistic architecture in which biological variability and competing model explanations remain visible through estimation, updating, and decision support. The objective is not to make digital twins maximally complex, but to prevent apparently precise model outputs from becoming stronger claims than the underlying biological and analytical evidence permits.

Keywords: Digital twin, Bioprocessing, Biological variability, State estimation, Uncertainty quantification, Hybrid modeling

Received: 17 June 2026; **Revised:** 18 September 2026; **Accepted:** 21 September 2026

INTRODUCTION

Digital twins are attractive to pharmaceutical biotechnology because bioprocess development repeatedly confronts a problem that conventional scale-up cannot eliminate: the state encountered at manufacturing scale is not simply a larger version of the state measured in a development vessel. Spatial gradients, physiological adaptation, changing cell populations, sensor limitations, and imperfect models can all intervene between laboratory observations and industrial performance. Recent analysis of AI-enabled bioprocess scale-up therefore places digital twins within a broader effort to connect mechanistic representation, data acquisition, and adaptive learning while acknowledging that scale-dependent uncertainty remains substantial [1]. The relevant question is

consequently not whether a model can predict a process variable, but whether the digital representation continues to describe the biological and engineering system for the decision being made.

The idea of a digital bioprocess replica already implies more than a standalone predictive algorithm. Computational models can represent unit operations, material transformations, operating conditions, and quality-relevant process relationships, but an integrated replica must also connect those representations to evolving process information [2]. For pharmaceutical biotechnology, this requirement becomes particularly stringent because the manufacturing system is simultaneously physical and biological. Cell physiology can change while bulk measurements remain apparently stable, model parameters can compensate for missing mechanisms, and an accurate short-horizon forecast can coexist with poor extrapolation beyond the state space used to train or calibrate it.

Increasing automation does not remove these distinctions. Reviews of AI-enabled bioprocess automation emphasize that multiscale biological complexity, safety requirements, and regulatory constraints make bounded human-machine collaboration more realistic than assuming that a self-driving platform can infer when its own biological representation has ceased to be valid [3]. The methodological position advanced here is therefore that biological uncertainty should be preserved inside the twin rather than hidden behind a single best estimate. The proposed approach separates observations from inferred states, parameter uncertainty from structural uncertainty, biological variability from unexplained residual error, and optimization from qualification. A digital twin becomes useful not when uncertainty disappears, but when the architecture retains enough information to show what is known, what is inferred, what may have changed, and which decisions remain defensible.

What a bioprocess digital twin must represent

A bioprocess twin first requires a trustworthy informational substrate. Autonomous biological laboratories increasingly depend on data infrastructures that preserve experimental context, metadata, provenance, and machine-readable relationships among measurements rather than treating the model as an isolated analytical object [4]. The same principle applies more strongly in manufacturing. A concentration value without sampling context, instrument status, process phase, calibration history, or transformation provenance may be numerically precise but insufficient to determine whether a model update is justified. Data lineage should therefore be considered part of twin validity, although it is not itself a biological state.

Model-based bioprocess development increasingly combines mechanistic models, statistical learning, process analytical technology, soft sensors, optimization, and digital-twin concepts [5]. These elements should not be collapsed into one computational block. The twin must distinguish at least the measured process state, the inferred biological state, the model representation used to connect them, and the decision layer that acts on the resulting information. Such separation allows disagreement to become diagnostically useful. A sensor residual can then indicate measurement deterioration, an unrepresented biological transition, parameter drift, or structural model failure rather than being automatically absorbed by whichever adaptive component is easiest to retune.

Mammalian-cell digital-twin concepts illustrate why these levels matter. Proposed architectures integrate in-line measurements with models of cell growth, metabolism, bioreactor operation, key performance indicators, and quality-relevant outputs [6]. Yet the existence of these connections does not establish that a change in a process parameter causes a particular critical quality attribute in every cell line, scale, or production phase. The proposed representation in this article therefore keeps cell state, reactor state, and product-relevant state analytically separate. Their connections are conditional relationships whose strength may depend on cell line, culture phase, operating history, analytical sensitivity, and scale. A trustworthy twin should preserve those conditions rather than replacing them with an undifferentiated “process state.”

Mechanistic models

Mechanistic models provide one way to constrain what the latent biological state can plausibly be. Genome-scale metabolic models for Chinese hamster ovary cells can organize metabolic capabilities, connect extracellular conditions to feasible intracellular flux patterns, and support interpretation of production phenotypes [7]. Their value lies in restricting inference using biological structure. Their limitation is equally important: stoichiometric feasibility does not encode every regulatory, signaling, epigenetic, secretory, stress-response, or population-level process that can influence manufacturing behavior. A mechanistic component can therefore narrow the state space without proving that the correct mechanism has been identified.

The same caution applies across scale. Computational scale-up frameworks coupling hydrodynamics, transport, and cellular responses show why well-mixed laboratory models can become inadequate when cells encounter

fluctuating oxygen, substrate, pH, or shear environments in larger reactors [8]. Exposure history can matter as much as nominal bulk conditions. For a digital twin, this means that a mechanistic state may need to include spatially or temporally integrated environmental exposure rather than only vessel-average measurements. Whether that additional resolution is necessary must remain process-specific; unnecessary mechanistic detail can create poorly identifiable models without improving the intended decision.

Hybrid modelling provides a practical response to incomplete mechanism. Mechanistic and data-driven models have complementary strengths in upstream pharmaceutical bioprocessing: mechanistic structure can enforce biological or physical plausibility, whereas statistical components can capture systematic patterns not represented explicitly in the governing equations [9]. The important methodological boundary is that residual learning should not be interpreted as proof that structural misspecification has disappeared. A data-driven correction can improve prediction while masking an incorrect mechanism.

Process-specific demonstrations show how the division of labor can work. In a fed-batch microalgal process, physics-based state reconciliation was combined with data-driven forecasting and optimization to support online simulation and future-state prediction [10]. This establishes the feasibility of assigning different computational components different functions. It does not establish that the same model architecture transfers to mammalian therapeutic-protein manufacture or that improved prediction confirms the underlying biology. In the proposed uncertainty-preserving twin, mechanistic constraints, learned residuals, and structural alternatives therefore remain separately identifiable.

Data-driven components

High-throughput technologies and automated workflows can accelerate design–build–test–learn cycles by generating larger and more structured experimental datasets [11]. For a digital twin, however, increased data volume is useful only when the measurements span states relevant to the intended manufacturing decisions. Dense observations from a narrow operating region can produce an apparently accurate model that fails during scale transition, equipment change, cell-line drift, or an uncommon disturbance.

Machine learning has demonstrated value across bioprocess development, including strain or cell-line development, process optimization, monitoring, prediction, and control, but practical deployment remains conditioned by data quality, model interpretation, and transferability [12]. A twin should therefore store the domain under which a data-driven component was trained and validated. Model identity should include the process phase, cell line or organism, sensor configuration, input range, prediction horizon, and version of the data pipeline from which the model learned.

AI-enabled biofoundries extend this logic by using active learning and automated experimentation to select informative experiments within iterative engineering workflows [13]. Such systems suggest a valuable role for the digital twin: uncertainty can guide what should be measured next. This is different from allowing an algorithm to certify its own predictions. The proposed architecture uses data-driven models to identify informative experiments or operating observations while reserving qualification as a separate evidentiary process.

Knowledge-oriented AI should also remain distinct from process sensing. Large language models, retrieval-augmented generation, and knowledge graphs may support literature synthesis, protocol interpretation, and engineering knowledge integration, but their usefulness depends on benchmarking, provenance, and controls against unsupported generation [14]. They should not be treated as equivalent to calibrated process analytical measurements. A digital twin may use such systems to contextualize evidence without promoting generated text to the status of direct process observation.

Time horizon introduces another form of domain dependence. Multistep-ahead modelling of mammalian cell culture has shown that predictive behavior changes with feature selection, algorithm choice, and forecasting horizon [15]. Consequently, a forecast should carry its temporal validity with it. A model that is reliable for near-term state estimation is not automatically a reliable predictor of later product-quality consequences. Preserving this distinction prevents a high-performing short-horizon model from acquiring unwarranted authority elsewhere in the manufacturing trajectory.

State estimation

Bioprocesses contain important variables that cannot be measured continuously, making state estimation central to digital-twin operation. Dynamic oxygen-uptake information, for example, can support soft-sensor estimation of viable-cell state in mammalian culture [16]. The output remains an inference conditioned on calibration, process phase, and the relationship between respiration and viable-cell behavior. The twin should therefore label the

variable as estimated rather than allowing the soft-sensor output to become indistinguishable from a direct measurement.

Interpretable machine-learning soft sensors provide another route to inferred state. In simulated industrial-scale penicillin fermentation, interpretable learners were compared for real-time estimation of product concentration, with performance varying among algorithms [17]. Such evidence supports the feasibility of interpretable state estimation but also demonstrates an important evidentiary boundary: performance against a process simulator is not equivalent to qualification against an operating pharmaceutical manufacturing process.

Mechanistic observers provide a different form of constraint. A soft sensor for protein-refolding processes combined intrinsic-fluorescence measurements with a mechanistic model so that the inferred trajectory remained consistent with the represented process dynamics [18]. Physical admissibility can make an estimate more useful, but a physically admissible estimate is not necessarily the uniquely correct biological explanation. The observation model, kinetic assumptions, and sensor informativeness remain potential uncertainty sources.

The distinction becomes sharper when parameters are estimated together with states. Regularized moving-horizon estimation has been applied to combined state and parameter estimation in bioprocess experimentation where partial observation and ill-conditioning complicate inference [19]. Parameter adaptation can reduce mismatch, but it should not be interpreted as evidence that the model structure is correct. A parameter may compensate for an omitted pathway, an unrepresented population shift, or a biased sensor.

Comparative work on closed-loop state and parameter estimation further shows that the preferred updating strategy depends on how uncertainty enters the underlying dynamics [20]. This supports a mechanism-discrimination approach rather than a universal updater. The proposed twin therefore asks not only “what state best fits the data?” but also “which source of mismatch would have to be true for this update to be scientifically defensible?”

That distinction is difficult to preserve in narrative form alone. **Table 1** therefore separates representative estimation situations according to what is actually observed, what is inferred, and what evidence would be required before a corrective action could be assigned to a biological rather than computational cause.

Table 1. Distinguishing corrective actions from competing explanations during bioprocess state estimation

Scientific configuration	Observable support	Mechanistic test and consequence
estimation situation / Latent interpretation being considered	Information directly available	Candidate control or model action / Recovery behavior that discriminates explanations / Decision consequence
Manuscript-derived estimation situation: Physiological proxy departs from expected culture trajectory Latent interpretation being considered: Change in viable-cell physiology	Calibrated process signal and recent operating history	Candidate control or model action: Update latent-state estimate while retaining sensor uncertainty Recovery behavior that discriminates explanations: Independent measurement confirms the same directional state change Decision consequence: State correction may inform control without redefining model structure
Manuscript-derived estimation situation: Soft-sensor prediction diverges while primary measurements remain coherent Latent interpretation being considered: Model-domain drift or unrepresented state	Process measurements plus prediction residual	Candidate control or model action: Suspend automatic reliance on the soft sensor and request corroboration Recovery behavior that discriminates explanations: Residual resolves after model-domain update but raw measurements remain internally consistent Decision consequence: Prediction component may require recalibration rather than process intervention
Manuscript-derived estimation situation: Mechanistic observer becomes inconsistent with one or more measurements Latent interpretation being considered: Sensor fault, parameter shift, or missing mechanism	Sensor observations and mechanistic constraints	Candidate control or model action: Test measurement integrity before widening parameter adaptation Recovery behavior that discriminates explanations: Independent sensing distinguishes measurement failure from persistent biological mismatch Decision consequence: Automatic parameter compensation is not yet justified

Manuscript-derived estimation situation: Parameter adaptation repeatedly restores short-term fit Latent interpretation being considered: Genuine kinetic change or structural compensation	Measurements, fitted states, updated parameter set	Candidate control or model action: Maintain parallel structural explanation rather than treating new parameter as truth Recovery behavior that discriminates explanations: New perturbation separates alternative models that previously fitted equally well Decision consequence: Structural uncertainty remains active
Manuscript-derived estimation situation: State and parameter estimates change together after a disturbance Latent interpretation being considered: Changed biological state with altered kinetics	Time-resolved measurements and estimator output	Candidate control or model action: Use bounded joint update Recovery behavior that discriminates explanations: Post-disturbance observations stabilize consistently across independent variables Decision consequence: Joint update may be actionable within the qualified disturbance domain
Manuscript-derived estimation situation: Residual remains systematic after plausible state and parameter updates Latent interpretation being considered: Structural inadequacy or out-of-domain operation	Measurements and persistent discrepancy	Candidate control or model action: Escalate from adaptation to model review Recovery behavior that discriminates explanations: Competing structure better explains new observations without implausible compensation Decision consequence: Twin should enter a restricted decision state

Note: The classification is a proposed methodological synthesis of the approved evidence. It is not a validated diagnostic scoring system, and the displayed actions require process-specific qualification.

State estimation also determines the architecture of the first conceptual figure. The key visual problem is not a linear path from sensor to model to controller, but a closed field in which estimated states can be revised while validation gates prevent an estimator from converting every discrepancy into a biological conclusion.




Figure 1. State estimation under competing biological, measurement, and model explanations

The figure represents state estimation as a closed-loop evidentiary field rather than a direct sensor-to-control pipeline. Process analytical measurements and contextual process information enter an estimation domain containing mechanistic constraints and data-driven estimators. Estimated biological and process states remain explicitly separated from direct observations. Parameter adaptation is permitted through a distinct update path, while persistent discordance can be routed to a structural-adequacy channel rather than being absorbed automatically by parameter fitting. Cross-scale and intended-use validation gates interrupt progression from estimation to control. Solid connectors represent relationships directly supported by the approved literature; dashed connectors denote the article's proposed diagnostic or escalation logic. Bidirectional connectors indicate iterative updating. No connector implies that predictive fit alone establishes a biological mechanism or product-quality consequence.

Biological variability

A digital twin that represents only a population-average cell state risks treating genuine biology as noise. Cellular variability can arise from cell age, division history, metabolic state, microenvironmental exposure, and other sources that generate functionally distinct subpopulations even under a common bulk process condition [21]. The manufacturing consequence is context dependent: heterogeneity may reduce consistency in one setting, have little detectable consequence in another, or represent adaptive behavior rather than failure. Its existence should therefore not be translated automatically into a negative quality interpretation.

The proposed twin treats biological variability as a distribution over plausible biological states whenever the evidence warrants that representation. This is different from simply increasing residual variance. Residual error describes disagreement between model and observation; biological heterogeneity proposes that different cells or subpopulations can genuinely occupy different states. The two explanations may coexist, and bulk measurements may not identify them uniquely.

Adaptive predictive control provides a methodological precedent for responding to changing biological dynamics rather than assuming a fixed model throughout operation [22]. Yet adaptation should occur because evidence supports a changed state or model relationship, not because every residual is presumed biological. A controller that continuously retunes itself to sensor bias or structural model error can become more numerically responsive while less biologically interpretable. Biological variability therefore requires a decision ledger that records what is known, what remains hidden, what action is permissible, and what observation should trigger reassessment.

Table 2. Decision states for preserving hidden biological variability inside a pharmaceutical bioprocess twin

Decision context	Evidence and limits	Permitted action	
Hidden biological state plausible / Disturbance or variability source	Information available at that state	Decision state / Permissible response	Observation that should trigger reassessment / Consequence for the digital twin
Hidden biological state that remains plausible: Emerging subpopulation or altered metabolic state Disturbance or variability source: Cell-state redistribution	Process conditions plus changing proxy	Decision state: Bulk variables remain stable while a physiological proxy shifts Permissible response: Increase biological surveillance before changing process set-points	Observation that should trigger reassessment: Independent biological measurement corroborates or contradicts the proxy Consequence for the digital twin: Broaden latent-state distribution rather than force a single-state estimate
Hidden biological state that remains plausible: Reversible adaptation across subpopulations Disturbance or variability source: Environmental fluctuation	Disturbance history and recovery measurements	Decision state: Process disturbance produces a transient heterogeneous response Permissible response: Maintain bounded adaptive control	Observation that should trigger reassessment: Recovery fails to return toward the prior biological distribution Consequence for the digital twin: Reclassify event from transient state change to possible model/domain change

Hidden biological state that remains plausible: Multiple phenotypes with similar mean output Disturbance or variability source: Cell age, division history, microenvironment	Population-average measurements plus heterogeneity-sensitive data	Decision state: Average culture behavior is preserved but variability increases Permissible response: Avoid interpreting the unchanged mean as unchanged biology	Observation that should trigger reassessment: Product-relevant or metabolic divergence emerges Consequence for the digital twin: Retain heterogeneity as an active state variable
Hidden biological state that remains plausible: Biological shift remains only one competing explanation Disturbance or variability source: Sensor drift, parameter change, missing mechanism, or biology	Sensor data and worsening model fit	Decision state: Residual error rises without independent evidence of biological change Permissible response: Diagnose competing causes before biological adaptation	Observation that should trigger reassessment: Orthogonal measurements localize the mismatch Consequence for the digital twin: Prevent residual error from being relabeled automatically as heterogeneity
Hidden biological state that remains plausible: Changed state distribution Disturbance or variability source: Sustained physiological adaptation	Convergent process and biological information	Decision state: Biological shift is corroborated across independent measurements Permissible response: Update biological-state representation within qualified bounds	Observation that should trigger reassessment: New state persists beyond the established adaptation window Consequence for the digital twin: Permit controlled model adaptation while retaining uncertainty
Hidden biological state that remains plausible: Uncharacterized biological condition Disturbance or variability source: Unseen perturbation or domain shift	Measurements inconsistent with established state repertoire	Decision state: Novel state lies outside the qualified biological domain Permissible response: Restrict automated decision authority	Observation that should trigger reassessment: Additional characterization establishes a new qualified domain Consequence for the digital twin: Escalate to model review and lifecycle requalification

Note: These decision states are a proposed synthesis, not empirically validated categories. Their purpose is to prevent bulk stability, prediction residuals, and biological heterogeneity from becoming interchangeable explanations.

Parameter and structural uncertainty

Uncertainty quantification is often added after prediction, yet automated biological systems require uncertainty to be tracked through preprocessing, data integration, modelling, experimental selection, and decision making [23]. For a pharmaceutical digital twin, uncertainty should therefore be represented by source. Measurement uncertainty describes uncertainty in what was observed; state uncertainty concerns the latent condition inferred from those observations; parameter uncertainty concerns unknown values within a chosen model; structural uncertainty concerns whether the chosen model family itself is adequate; and domain uncertainty concerns whether the current process still belongs to the region in which those claims were qualified.

Predictive uncertainty is not equivalent to structural adequacy. Comparative work using conformal prediction and Gaussian-process models in chromatography shows that uncertainty estimates and calibration behavior can differ materially among modelling approaches and that mechanistic information can alter predictive uncertainty [24]. Such intervals are valuable for describing prediction reliability within a domain. They do not demonstrate that the underlying mechanistic interpretation is correct. A calibrated model can remain structurally wrong while producing useful predictions.

Parameter uncertainty, by contrast, can often be incorporated directly into control design. Robust tube-based nonlinear model predictive control of mammalian cell culture has demonstrated computational strategies in which uncertain kinetic parameters generate bounded trajectories or safety corridors rather than a single nominal future [25]. This supports an important design principle for digital twins: a control decision can be conditioned on a range of plausible process behaviors. The evidence remains computational, however, and cannot be converted into a general claim that such corridors are prospectively qualified for manufacturing or regulatory use.

Probabilistic data-driven control extends the same concept. Gaussian-process state-space models can propagate state-dependent predictive uncertainty and incorporate probabilistic back-offs into model predictive control [26].

This provides a methodological alternative to making decisions from point estimates alone. Its guarantees remain conditional on modelling and calibration assumptions, and uncertainty outside the represented domain may not be captured correctly.

The proposed digital twin therefore preserves uncertainty as a set of interacting but non-interchangeable channels. Parameter distributions may be updated recursively when the model structure remains credible. Structural alternatives require a higher evidentiary threshold because improving fit through parameter adaptation can conceal missing biology. Predictive intervals can bound decisions without becoming mechanistic explanations. Biological-state distributions can widen when evidence supports heterogeneity without being replaced by generic residual variance. When these channels disagree, the appropriate response is not necessarily to compute a narrower estimate. It may be to restrict decision authority, obtain discriminating measurements, or return the twin to an earlier level of model qualification.

This distinction is central to the article's methodological claim. A pharmaceutical biotechnology digital twin should not be defined by how confidently it can output one future trajectory. It should be defined by whether it can preserve the scientifically important alternatives that remain plausible while the process evolves.

Updating the twin during production

A production digital twin must remain synchronized with the physical process without assuming that every new observation warrants the same kind of update. A cyber-physical implementation for continuous clarification illustrates the practical importance of connecting sensors, equipment, models, and computational decisions through an explicit physical-digital interface [27]. Such synchronization is necessary but insufficient: updating a numerical state does not establish that the underlying biological model remains valid.

Biological change may require a qualitatively different response. Hybrid modelling combined with transfer learning has been used to adapt bioprocess models to a different biological production strain using limited target-domain data [28]. This supports distinguishing ordinary state updates from parameter adaptation and domain transfer. The proposed twin therefore records which layer has changed. A state update assumes the representation remains adequate; a parameter update modifies quantities within that representation; domain transfer acknowledges that relationships learned previously may no longer be directly transportable.

Model validity can also decay after deployment. An MLOps-oriented soft-sensor study demonstrated that process deviations can generate concept drift severe enough to degrade predictive performance and motivate retraining [29]. Retraining, however, is not equivalent to requalification. A changed model may restore predictive fit while altering its operating domain, uncertainty calibration, or relationship to the decisions it informs. Model-health status should consequently remain an explicit component of the twin.

Decision support under uncertainty

The output of a digital twin becomes consequential when it changes manufacturing action. Reinforcement-learning approaches demonstrate that control policies can adapt under stochasticity and plant-model mismatch [30]. Such results establish computational feasibility, not automatic suitability for pharmaceutical control. An optimized policy can still be unacceptable if its recommendation depends on an unqualified biological state, poorly calibrated uncertainty, or an extrapolation beyond the model's intended-use domain.

Likewise, robust multi-setpoint and trajectory tracking has been demonstrated computationally for uncertain bioprocess systems [31]. Tracking accuracy is nevertheless a different endpoint from product-quality assurance. A controller can follow its modeled objective successfully while the relationship between that objective and a critical quality attribute remains uncertain. Decision support therefore requires an explicit boundary between computationally preferred action and scientifically permissible action.

Table 3 organizes failure conditions beginning with the misleading pattern rather than with the modeling technique. This exception-first structure is needed because similar numerical outputs can require very different responses depending on whether the dominant uncertainty originates in sensing, biology, model structure, scale transfer, or lifecycle degradation.

Table 3. Exception-first triage of uncertainty before digital-twin recommendations become manufacturing actions

Observed configuration	Mechanism and competing diagnosis	Corrective discrimination
------------------------	-----------------------------------	---------------------------

Observable signature	Misleading or failure pattern / Likely origin evaluation / Competing diagnosis	Corrective qualification test / Residual consequence after mitigation
Action sensitivity rises while estimate confidence deteriorates	Misleading or failure pattern: Optimizer recommends a large control change despite widening state uncertainty Likely origin requiring evaluation: Decision logic acting on an unstable inferred state Competing diagnosis: Genuine rapid biological transition	Corrective qualification test: Re-estimate state using independent observations and bounded uncertainty Residual consequence after mitigation: Automated action remains restricted until the state is discriminated
Residual decreases but parameter values move outside previously credible behavior	Misleading or failure pattern: Model fit improves after aggressive parameter adaptation Likely origin requiring evaluation: Parameter compensation Competing diagnosis: Missing or changed model structure	Corrective qualification test: Test competing structures against a discriminating perturbation Residual consequence after mitigation: Restored fit alone cannot recover structural confidence
Short-horizon agreement with accumulating spatial or physiological mismatch	Misleading or failure pattern: Predictor remains accurate locally after scale transition Likely origin requiring evaluation: Unrepresented scale dependence Competing diagnosis: Temporary benign disturbance	Corrective qualification test: Validate relevant exposure histories and scale-sensitive variables Residual consequence after mitigation: Decision use is limited to the verified transfer domain
Set-point tracking succeeds without corroborating quality evidence	Misleading or failure pattern: Controller tracks its computational target while product relevance remains uncertain Likely origin requiring evaluation: Objective–quality disconnect Competing diagnosis: Objective is genuinely adequate but incompletely measured	Corrective qualification test: Test the objective against product-relevant measurements Residual consequence after mitigation: Tracking performance cannot substitute for quality qualification
Error falls after retraining but learned relationships or calibration change	Misleading or failure pattern: Retrained model restores prediction after drift Likely origin requiring evaluation: Lifecycle model change Competing diagnosis: Temporary data-quality event	Corrective qualification test: Reassess model domain, calibration, interfaces, and intended use Residual consequence after mitigation: Retrained version requires controlled requalification
Similar historical fit but divergent future trajectories	Misleading or failure pattern: Multiple plausible models recommend different actions Likely origin requiring evaluation: Structural uncertainty Competing diagnosis: Insufficiently informative data	Corrective qualification test: Acquire observations chosen to discriminate the competing explanations Residual consequence after mitigation: Decision authority should remain bounded while alternatives persist

Note: The failure modes and triage logic constitute a proposed synthesis. They do not define universal regulatory thresholds, validation criteria, or automated release rules.

The same boundary is the organizing principle of **Figure 2**: cellular and process-state estimates may inform a product-quality interpretation, but they cannot cross automatically into a manufacturing decision when the relationship remains uncertain.




Figure 2. Preventing uncertain cell and process states from becoming unjustified product-quality decisions

The multilevel map separates biological state, process state, product-relevant interpretation, and manufacturing action. Evidence-supported connections are shown independently from the article's proposed diagnostic and escalation links. A central decision boundary blocks automatic translation of a predicted cell or process change into a product-quality conclusion when measurement validity, structural adequacy, scale transfer, or quality linkage remains unresolved. Lower uncertainty channels return discordant evidence to the appropriate earlier analytical level rather than allowing an optimizer to absorb every discrepancy. Solid connectors denote evidence-supported relationships; dashed connectors denote proposed conditional or escalation logic; bidirectional connectors represent recursive updating.

A probabilistic digital-twin architecture

Cross-scale biological control increasingly links molecular regulation with macroscopic process behavior, reinforcing that distinct biological and engineering layers can be coupled without being treated as identical [32]. On that basis, this article proposes an uncertainty-preserving digital-twin architecture composed of six interacting representations: direct observations; latent biological and process states; parameter distributions; competing structural-model explanations; model-health and domain status; and bounded decision constraints.

The architecture is probabilistic because a single best state is not always scientifically defensible. It is also hierarchical because uncertainty should return to the level where it originates. Measurement disagreement should first challenge the observation layer; biological discordance should broaden or revise biological-state hypotheses; repeated parameter compensation should trigger structural scrutiny; and out-of-domain behavior should restrict decision authority. The purpose is not to maintain every conceivable model simultaneously, but to preserve alternatives that remain materially relevant to the manufacturing decision.

Qualification and lifecycle maintenance

End-to-end digital-twin software for continuous mRNA manufacturing demonstrates how mechanistic models of multiple unit operations can be assembled within a common computational environment and interrogated through sensitivity analysis [33]. Such integration is an important engineering step, but it does not by itself constitute prospective closed-loop manufacturing qualification.

Qualification should therefore be modular and lifecycle-based. Measurement interfaces require analytical qualification; state estimators require domain-specific validation; parameter updates require identifiability and plausibility checks; structural changes require renewed model-adequacy assessment; transferred or retrained components require verification in the target domain; and decision policies require evidence that their

optimization objective remains relevant to the intended pharmaceutical endpoint. End-to-end testing is then necessary to examine propagation across interfaces rather than replacing these component-level checks.

Under this view, lifecycle maintenance is not merely software maintenance. A new sensor, cell line, scale, operating strategy, model version, retraining event, or newly observed biological state can change what the twin is justified in claiming. The appropriate response may be recalibration, parameter re-estimation, model revision, restricted decision authority, additional experimentation, or formal requalification. A digital twin remains credible only while its evidentiary boundaries evolve as deliberately as its algorithms.

CONCLUSION

Pharmaceutical-biotechnology digital twins should preserve biological uncertainty rather than convert it prematurely into a single deterministic state. Measurements, inferred biological states, parameters, structural assumptions, domain validity, and decisions represent different evidentiary levels and should remain distinguishable during operation.

The proposed architecture treats uncertainty as information that guides estimation, experimentation, updating, and limits on automation. Its central safeguard is an explicit boundary between a computationally available prediction and a scientifically justified manufacturing action. This framework is methodological rather than empirically validated, and its implementation must remain product-, process-, scale-, sensor-, and intended-use specific. The value of a digital twin should therefore be judged not only by predictive precision, but also by how faithfully it communicates when biological alternatives remain unresolved.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Yuan S, Xu V, Muddana C, Sureshkumar P, Tang YJ. From design-build-test-learn cycles to AI-driven digital twins for bioprocess scale-up in the Genesis Mission era. *Curr Opin Biotechnol.* 2026;100:103516. doi:10.1016/j.copbio.2026.103516
2. Smiatek J, Jung A, Bluhmki E. Towards a digital bioprocess replica: Computational approaches in biopharmaceutical development and manufacturing. *Trends Biotechnol.* 2020;38(10):1141-53. doi:10.1016/j.tibtech.2020.05.008
3. Helleckes LM, Putz S, Gupta K, Franzreb M, Garcia Martin H. Perspectives for artificial intelligence in bioprocess automation. *Curr Opin Biotechnol.* 2026;97:103392. doi:10.1016/j.copbio.2025.103392
4. Sommer LM, Groves T, Santos A. Beyond the model: Data infrastructure as the foundation for autonomous virtual laboratories. *Curr Opin Biotechnol.* 2026;97:103434. doi:10.1016/j.copbio.2025.103434
5. Mu'azzam K, da Silva FVS, Murtagh J, Sousa Gallagher MJ. A roadmap for model-based bioprocess development. *Biotechnol Adv.* 2024;73:108378. doi:10.1016/j.biotechadv.2024.108378
6. Park SY, Park CH, Choi DH, Hong JK, Lee DY. Bioprocess digital twins of mammalian cell culture for advanced biomanufacturing. *Curr Opin Chem Eng.* 2021;33:100702. doi:10.1016/j.coche.2021.100702
7. Park SY, Choi DH, Song J, Lakshmanan M, Richelle A, Yoon S, et al. Driving towards digital biomanufacturing by CHO genome-scale models. *Trends Biotechnol.* 2024;42(9):1192-203. doi:10.1016/j.tibtech.2024.03.001
8. Wang G, Haringa C, Noorman H, Chu J, Zhuang Y. Developing a computational framework to advance bioprocess scale-up. *Trends Biotechnol.* 2020;38(8):846-56. doi:10.1016/j.tibtech.2020.01.009
9. Tsopanoglou A, Jiménez del Val I. Moving towards an era of hybrid modelling: Advantages and challenges of coupling mechanistic and data-driven models for upstream pharmaceutical bioprocesses. *Curr Opin Chem Eng.* 2021;32:100691. doi:10.1016/j.coche.2021.100691

10. Zhang D, Del Rio-Chanona EA, Petsagkourakis P, Wagner J. Hybrid physics-based and data-driven modeling for bioprocess online simulation and optimization. *Biotechnol Bioeng.* 2019;116(11):2919-930. doi:10.1002/bit.27120
11. Petzold CJ, Mukhopadhyay A. From bench to biofactory: High-throughput technologies and automated workflows to accelerate biomanufacturing. *Curr Opin Biotechnol.* 2025;94:103320. doi:10.1016/j.copbio.2025.103320
12. Helleckes LM, Hemmerich J, Wiechert W, von Lieres E, Grünberger A. Machine learning in bioprocess development: From promise to practice. *Trends Biotechnol.* 2023;41(6):817-35. doi:10.1016/j.tibtech.2022.10.010
13. Chen J, Singh N, Lu J, Lane ST, Zhao H. Artificial intelligence-powered biofoundries for protein engineering and metabolic engineering. *Curr Opin Biotechnol.* 2025;96:103380. doi:10.1016/j.copbio.2025.103380
14. Li W, Mao Z, Xiao Z, Liao X, Koffas M, Chen Y, et al. Large language model for knowledge synthesis and AI-enhanced biomanufacturing. *Trends Biotechnol.* 2025;43(8):1864-75. doi:10.1016/j.tibtech.2025.02.008
15. Park SY, Kim SJ, Park CH, Kim J, Lee DY. Data-driven prediction models for forecasting multistep ahead profiles of mammalian cell culture toward bioprocess digital twins. *Biotechnol Bioeng.* 2023;120(9):2494-508. doi:10.1002/bit.28405
16. Winter M, Achleitner L, Satzer P. Soft sensor for viable cell counting by measuring dynamic oxygen uptake rate. *N Biotechnol.* 2024;83:16-25. doi:10.1016/j.nbt.2024.06.001
17. Acosta-Pavas JC, Robles-Rodriguez CE, Griol D, Daboussi F, Aceves-Lara CA, Corrales DC. Soft sensors based on interpretable learners for industrial-scale fed-batch fermentation: Learning from simulations. *Comput Chem Eng.* 2024;187:108736. doi:10.1016/j.compchemeng.2024.108736
18. Igwe CL, Gisperm F, Kierein M, Brichtová EP, Spadiut O, Müller DF. Mechanistic soft-sensor design for protein refolding processes based on intrinsic fluorescence measurements. *Comput Chem Eng.* 2024;187:108734. doi:10.1016/j.compchemeng.2024.108734
19. Tuveri A, Nakama CSM, Matias J, Holck HE, Jäschke J, Imsland L, et al. A regularized moving horizon estimator for combined state and parameter estimation in a bioprocess experimental application. *Comput Chem Eng.* 2023;172:108183. doi:10.1016/j.compchemeng.2023.108183
20. Matias J, Swartz CLE. State and parameter estimation in closed-loop dynamic real-time optimization—a comparative study. *Comput Chem Eng.* 2025;194:108932. doi:10.1016/j.compchemeng.2024.108932
21. Eigenfeld M, Schwaminger SP. Cellular variability as a driver for bioprocess innovation and optimization. *Biotechnol Adv.* 2025;79:108528. doi:10.1016/j.biotechadv.2025.108528
22. Jabarivelisdeh B, Carius L, Findeisen R, Waldherr S. Adaptive predictive control of bioprocesses with constraint-based modeling and estimation. *Comput Chem Eng.* 2020;135:106744. doi:10.1016/j.compchemeng.2020.106744
23. Wiechert W, Helleckes LM, Nöh K. Building trust in automated experimentation: Uncertainty quantification in the era of high-throughput biolabs. *Curr Opin Biotechnol.* 2025;96:103382. doi:10.1016/j.copbio.2025.103382
24. Pham TD, Bassett R, Aickelin U. Capturing uncertainty in black-box chromatography modelling using conformal prediction and Gaussian processes. *Comput Chem Eng.* 2025;199:109136. doi:10.1016/j.compchemeng.2025.109136
25. Dewasme L, Mäkinen M, Chotteau V. Multivariable robust tube-based nonlinear model predictive control of mammalian cell cultures. *Comput Chem Eng.* 2024;183:108592. doi:10.1016/j.compchemeng.2024.108592
26. Bradford E, Imsland L, Zhang D, del Rio Chanona EA. Stochastic data-driven model predictive control using Gaussian processes. *Comput Chem Eng.* 2020;139:106844. doi:10.1016/j.compchemeng.2020.106844
27. Banerjee S, Jesubalan NG, Kulkarni A, Agarwal A, Rathore AS. Developing cyber-physical system and digital twin for smart manufacturing: Methodology and case study of continuous clarification. *J Ind Inf Integr.* 2024;38:100577. doi:10.1016/j.jii.2024.100577
28. Riezzo L, Kay H, Feng Y, Jing K, Zhang D. Accelerating bioprocess digital twin development by integrating hybrid modelling with transfer learning. *Chem Eng J.* 2025;511:162018. doi:10.1016/j.cej.2025.162018
29. Metcalfe B, Acosta-Pavas JC, Robles-Rodriguez CE, Georgakilas GK, Dalamagas T, Aceves-Lara CA, et al. Towards a machine learning operations (MLOps) soft sensor for real-time predictions in industrial-scale fed-batch fermentation. *Comput Chem Eng.* 2025;194:108991. doi:10.1016/j.compchemeng.2024.108991
30. Petsagkourakis P, Sandoval IO, Bradford E, Zhang D, del Rio-Chanona EA. Reinforcement learning for batch bioprocess optimization. *Comput Chem Eng.* 2020;133:106649. doi:10.1016/j.compchemeng.2019.106649

31. Espinel-Ríos S, Avalos JL, del Rio Chanona EA, Zhang D. Reinforcement learning for efficient and robust multi-setpoint and multi-trajectory tracking in bioprocesses. *Comput Chem Eng.* 2025;202:109297. doi:10.1016/j.compchemeng.2025.109297
32. Treinen C, Peternell C, Noll P, Magosch O, Hausmann R, Henkel M. Molecular process control for industrial biotechnology. *Trends Biotechnol.* 2025;43(9):2117-32. doi:10.1016/j.tibtech.2025.04.003
33. Shahab MA, Stover NM, Al-Mahayni H, Srisuma P, Wang J, Shin S, et al. End-to-end digital twin software for continuous mRNA manufacturing. *Int J Pharm.* 2025;686:126301. doi:10.1016/j.ijpharm.2025.126301