



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

AI-Guided Bioprocess Control Should Optimize Critical Quality Attributes Before Productivity

Sven Larsson^{1*}, Erik Johansson², Anna Nilsson³, Lars Andersson⁴

¹Department of AI-Guided Bioprocess Control, Faculty of Pharmacy, Swedish University of Agricultural Sciences, Uppsala, Sweden.

²Department of Critical Quality Attributes and Process Optimization, Faculty of Pharmacy, KTH Royal Institute of Technology, Stockholm, Sweden.

³Department of Bioprocess AI and Productivity, Faculty of Pharmaceutical Sciences, Lund University, Lund, Sweden.

⁴Department of Quality-Driven Biomanufacturing, Faculty of Pharmacy, University of Gothenburg, Gothenburg, Sweden.

*Email: sven.larsson@slu.se

ABSTRACT

Artificial intelligence is increasingly capable of extracting process state from multivariate bioprocess data, forecasting trajectories, and proposing or executing control actions. Yet an optimization system can be technically sophisticated while pursuing the wrong objective. In biopharmaceutical manufacturing, maximizing titer, throughput, or equipment utilization does not by itself establish preservation of product quality. This Systems Perspective argues that AI-enabled control should therefore treat critical quality attributes as primary constraints or decision targets and productivity as an objective optimized within an acceptable quality-relevant operating domain. The distinction requires more than adding quality measurements to an existing productivity loop. Observed sensor signals, inferred biological states, predicted quality attributes, controller outputs, and validated product evidence have different evidentiary status and should remain separable. Multivariate process analytical technology, soft sensors, physics-informed and transfer-learning models, and adaptive control can shorten the path from observation to intervention, but each introduces uncertainty, transferability limits, and qualification requirements. We develop a quality-first control logic in which increasingly consequential actions require progressively stronger evidence that the relevant product-quality state remains within the validated domain. The resulting architecture is proposed as a systems-level manufacturing principle rather than an empirically validated universal control strategy.

Keywords: Bioprocess control, Critical quality attributes, Process analytical technology, Soft sensors, Machine learning, Biopharmaceutical manufacturing

Received: 08 June 2026; **Revised:** 14 September 2026; **Accepted:** 16 September 2026

INTRODUCTION

Artificial intelligence and machine learning are moving from retrospective bioprocess analysis toward process monitoring, prediction, optimization, and control, while analytical platforms are simultaneously becoming faster and more capable of resolving product-relevant information [1-3]. These developments make automated decision loops increasingly plausible. They do not, however, determine what such loops should optimize. A controller that accurately maximizes a measurable productivity endpoint can still be misaligned with pharmaceutical

manufacturing if the manipulated state alters product attributes that are measured incompletely, inferred indirectly, or detected only after the action has become difficult to reverse.

The central issue is therefore not whether AI can optimize a bioprocess, but which evidence should govern optimization authority. This Perspective proposes that critical quality attributes (CQAs) should define the primary decision boundary, whereas productivity should be optimized within a quality-compatible operating region. The proposal deliberately separates five stages that are often compressed into the single phrase “AI control”: observation, state inference, quality inference, action selection, and post-action verification. A process measurement is not automatically a biological-state measurement; a predicted CQA is not automatically equivalent to a qualified analytical result; and a successful model fit does not establish that the resulting control action is transferable to another cell line, product, scale, sensor configuration, or manufacturing regime.

Why productivity-first optimization is incomplete

Productivity is attractive as an optimization objective because it is operationally important and often comparatively easy to quantify. Biopharmaceutical quality, by contrast, is multidimensional. Antibody glycosylation illustrates the problem: glycan distributions can respond to cell biology, media composition, culture conditions, and downstream strategies, while their relevance depends on the product and intended function [4]. Thus, a process state that produces more material is not necessarily informationally sufficient to establish that the material remains equivalent in quality.

Machine learning can identify cultivation settings associated with improved CHO-process outcomes and can support systematic optimization of variables that would be difficult to explore exhaustively [5]. Such capability should not be confused with proof that the chosen objective function represents the full pharmaceutical objective. If model training emphasizes titer or another productivity metric while quality information is delayed or absent, the optimizer can preferentially learn the portion of process space that is well observed rather than the portion most relevant to product acceptability.

The same distinction persists in intensified and continuous manufacture. Pilot-scale end-to-end continuous monoclonal-antibody production demonstrates that upstream and downstream operations can be technically integrated [6]. Broader continuous-processing experience likewise identifies monitoring, process understanding, equipment integration, and control as enabling conditions for realizing consistent performance [7]. Continuous operation therefore strengthens rather than removes the need for quality-aware control. Faster material flow can shorten the interval available to recognize a quality-relevant deviation, making the informational architecture of the control loop at least as important as its productivity objective.

A quality-first hierarchy does not imply that productivity and quality are intrinsically antagonistic. In many operating regions they may improve together. The proposed distinction is narrower: productivity should not be granted decision priority merely because it is easier to measure. Where quality and productivity are concordant, quality constraints should be inexpensive. Where they diverge, the architecture should make that divergence visible before productivity optimization acquires actuator authority.

Critical quality attributes as control targets

A CQA-centered controller requires more than inserting a final-product assay into a conventional feedback loop. The relevant quality state must be observed or inferred at a latency compatible with the process decision. Attribute-focused process analytical technology can automate sampling and analytical preparation, reducing the delay between manufacturing state and product-quality information [8]. This creates the possibility of directing control toward attributes closer to the pharmaceutical endpoint rather than relying exclusively on upstream process parameters as proxies.

The alternative is model-mediated CQA estimation. Raman spectroscopy combined with convolutional neural networks has been used to predict cell-culture product-quality variables across differing process runs and cell-line contexts [9]. Such results are important because they move sensing closer to the desired control target. They also establish an evidentiary boundary: a predicted quality value remains conditional on the training domain, spectral quality, reference assay, model architecture, and operating regime that supported its development.

For quality-first control, CQAs can therefore occupy three different positions. Some may be measured directly with sufficient latency for process intervention. Others may be estimated through a validated surrogate or soft sensor. Still others may remain effectively delayed endpoints and can only constrain control indirectly through qualified intermediate states. These categories should not be collapsed. A controller should know whether its

“quality signal” is an analytical observation, a model estimate, or a downstream outcome that has not yet been observed.

The proposed decision principle is consequently quality relevance before optimization convenience. A variable becomes valuable to AI control not merely because it is available continuously, but because its relationship to a quality-relevant state is sufficiently understood for the intended action. This preserves a distinction between process visibility and pharmaceutical decision utility.

Multivariate process measurements

Bioprocesses rarely provide one measurement that uniquely represents the state relevant to quality. Inline Raman spectroscopy can support monitoring and control in high-cell-density perfusion, while spectral performance remains sensitive to the measurement environment and process context. Imaging-based approaches can expose morphological transitions in cell-therapy manufacturing, and broader PAT architectures combine physical sensors, spectroscopy, analytical assays, and data-driven estimators [10-12]. These modalities are complementary precisely because they observe different aspects of the process.

The quality-first problem is therefore one of state discrimination, not sensor accumulation. A glucose estimate, a morphology distribution, an automated attribute assay, and a conventional process measurement differ in specificity, latency, biological proximity, and dependence on model inference. An AI system that simply concatenates them into a feature vector risks hiding these differences inside the prediction model.

Instead, multivariate measurements should retain provenance. Direct observations should remain distinguishable from transformed measurements and model-derived states. Discordant signals should remain visible rather than automatically averaged away. Most importantly, the path from a sensor pattern to a CQA-relevant interpretation should contain an explicit discrimination step: what competing biological or analytical explanations could generate the same observed pattern, and what additional observation would distinguish them?

This is the role assigned to **Table 1**. Rather than ranking sensors, it separates the evidence available at a control decision from the hidden state that the controller would otherwise be tempted to assume.

Table 1. Distinguishing observed process evidence from CQA-relevant hidden states in multivariate control decisions

Decision context	Evidence and limits	Permitted action
control state / Hidden state inference / Quality-relevant interpretation	What is directly observed	Observation needed to discriminate alternatives / Consequence for AI control
Manuscript-derived control state: Spectroscopic change during high-density culture Hidden state requiring inference: Whether the change reflects metabolite dynamics, optical interference, or a broader physiological transition Quality-relevant interpretation: Quality relevance remains conditional until the spectral change is linked to the product state	Shift in multivariate spectral signal and associated process variables	Observation needed to discriminate alternatives: Orthogonal process or analytical measurement targeted to the competing explanations Consequence for AI control: Continue estimation only within the qualified spectral domain; escalate when signal origin is unresolved
Manuscript-derived control state: Stable conventional CPPs with changing multivariate fingerprint Hidden state requiring inference: Emerging cellular or metabolic state not visible in conventional variables Quality-relevant interpretation: Nominal CPPs do not by themselves establish a stable quality trajectory	Temperature, pH, gases, feed variables remain nominal while a higher-dimensional signal changes	Observation needed to discriminate alternatives: Measurement capable of resolving the suspected cellular, metabolic, or product-attribute change Consequence for AI control: Prevent “all CPPs in range” from being treated as sufficient evidence for unrestricted optimization
Manuscript-derived control state: Morphological population shift Hidden state requiring inference: Whether the population shift represents desired progression, stress, process perturbation, or donor-dependent behavior Quality-relevant interpretation: Morphology can support process-state discrimination but is not automatically a release attribute	Image-derived or latent morphology distribution changes	Observation needed to discriminate alternatives: Qualified functional or product-relevant measurement appropriate to the process Consequence for AI control: Use morphology as state evidence, not as an automatic substitute for potency or quality testing

Manuscript-derived control state: Automated product-attribute result becomes available Hidden state requiring inference: Whether the sampled result is representative of the material affected by the proposed intervention Quality-relevant interpretation: Higher pharmaceutical relevance than a remote process proxy, but still assay- and sampling-dependent	Near-real-time analytical result linked more directly to product state	Observation needed to discriminate alternatives: Confirmation of assay performance, sampling representativeness, and temporal alignment Consequence for AI control: Permit greater control authority only when analytical validity and process alignment are established
Manuscript-derived control state: Conflicting sensor modalities Hidden state requiring inference: Sensor fault, model failure, spatial heterogeneity, or genuine biological transition Quality-relevant interpretation: Discordance itself becomes decision-relevant evidence	Different measurement streams support different process interpretations	Observation needed to discriminate alternatives: Independent measurement selected to discriminate technical from biological causes Consequence for AI control: Reduce automation authority until the conflict is resolved

Note: The decision states and control consequences are proposed synthesis derived from the manuscript evidence. They do not define numerical alarm limits or universal CQA thresholds.

The same logic can be represented spatially. **Figure 1** makes the central boundary explicit: multivariate sensing can narrow uncertainty about the process, but it does not acquire direct actuator authority until observation, inference, and CQA relevance have been distinguished.




Figure 1. Evidence gates between multivariate sensing and CQA-relevant control authority

Soft sensors

Soft sensors extend process visibility when an important variable cannot be measured continuously or economically by a direct analytical method. Automated machine-learning workflows can use routinely available process information to estimate amino-acid concentrations in mammalian perfusion cultures, although performance can vary across analytes and can improve when additional informative measurements are available [13]. The relevant point for quality-first control is that a soft sensor creates an estimate, not a new physical measurement.

Calibration strategy therefore becomes part of the control architecture. Global Raman regression models developed across different cell lines have reduced the amount of new calibration data required when transferring

prediction capability to another cell-culture context [14]. By contrast, just-in-time approaches can select locally relevant observations and adapt prediction to the neighborhood of the current process state; a variational-autoencoder-based JITL approach improved prediction relative to tested comparator models across the reported cell-culture runs [15].

These strategies solve different problems. Global calibration attempts to make a model transportable across related domains, whereas local adaptation attempts to remain useful when the current state differs from the global training distribution. Neither strategy establishes perpetual validity. The former may fail when source and target domains are insufficiently related; the latter may become unreliable when the local neighborhood itself is sparse or misleading.

A quality-first architecture should consequently retain a distinction between measurement confidence and model confidence. A sensor may be analytically stable while its mapping to the inferred variable drifts, or a model may be statistically stable while the underlying measurement becomes unreliable. Combining both into one opaque confidence score would obscure the failure mode that the control system needs to diagnose.

Predictive models

Predictive bioprocess models face a recurrent asymmetry: industrial decisions may require predictions in regions where direct high-quality manufacturing data are scarce. Different modeling strategies address this limitation in different ways. Physics-informed neural networks constrain learning with process structure; multi-fidelity methods combine evidence of unequal cost or accuracy; and transfer learning reuses information from a source domain while recognizing that poorly matched domains can produce negative transfer [16-18]. These approaches are therefore better understood as strategies for managing limited information than as interchangeable routes to a universally generalizable model.

Synthetic process trajectories offer another way to supplement sparse historical data. In small-data batch-monitoring settings, generated in-silico trajectories can provide additional information for constructing multivariate monitoring models [19]. Their usefulness remains conditional on the fidelity of the generating model. Synthetic data can expand a learning set; they cannot independently demonstrate that a controller will behave correctly in a manufacturing state the generator itself represents poorly.

Industrial monitoring can also exploit structural similarity across products. Deep multistage monitoring has been developed for real-time surveillance of biopharmaceutical processes where no product-specific history exists, using information from other manufacturing recipes rather than waiting for a large new-product history to accumulate [20]. This is a powerful direction, but its governing question is again transportability: which relationships are shared strongly enough to transfer, and which are product-, cell-line-, site-, scale-, or sensor-specific?

The quality-first interpretation is that predictive performance must be attached to a domain statement. “Validated model” is incomplete unless the relevant operating region, product context, measurement configuration, and process state are also specified. A high-performing model inside one domain should lose decision authority when the process moves into a state that has not been adequately represented. Generalization is therefore not a permanent property of a model class; it is a conditional relationship among the model, the evidence used to build it, and the manufacturing state in which it is asked to act.

Closed-loop control

Closed-loop control turns inferential error into physical consequence. Adaptive model-predictive-control methods can combine learned process representations, online parameter estimation, constraints, and receding-horizon optimization, demonstrating how an AI-enabled controller can adapt decisions as a modeled fed-batch process evolves [21]. Such studies establish technical plausibility, not GMP validation. Their relevance to pharmaceutical manufacturing depends critically on the objective and constraints supplied to the controller.

The objective-function problem is visible beyond upstream culture. Machine-learning surrogate models can accelerate optimization of multistep ion-exchange chromatography by approximating expensive process calculations and searching a constrained design space [22]. What the optimizer regards as success is determined by the outcomes and constraints encoded in that design problem. A fast optimizer cannot compensate for an incomplete pharmaceutical objective.

Digital bioprocess replicas offer a systems-level mechanism for connecting measurements, models, process understanding, and prospective control calculations [23]. PAT can further shorten the observation-to-action interval, but implementation depends on measurement reliability, analytical treatment, and the maturity of the

sensing technology in the relevant unit operation [24]. Integrated continuous manufacturing adds an additional boundary: regulatory and validation questions attach to the coordinated system, not simply to the performance of isolated models or unit operations [25].

The proposed quality-first controller therefore uses a constrained authority hierarchy. A productivity-improving action is permissible only when the measurements and model state required to judge its quality relevance remain inside their qualified domains. When a direct quality measurement is unavailable, the controller may rely on a validated surrogate, but the uncertainty attached to that surrogate should narrow rather than expand automated authority. Where model evidence becomes discordant, out-of-domain, or insufficiently linked to the relevant CQA, the proper response is not necessarily to stop the process; it is to reduce the scope of automatic optimization and require a different evidence path.

Table 2 translates this principle across manufacturing scales and model roles. Its purpose is to prevent successful control at one analytical level from being interpreted automatically as evidence of transferability at the next.

Table 2. Cross-scale qualification boundaries for quality-first closed-loop bioprocess control

Source state and scale		Qualification endpoint
Manufacturing control case / Process or cellular state represented	Measurement or model evidence available / Hidden quality state and scale sensitivity	Qualification test before broader authority / Permissible consequence and unresolved question
Manufacturing control case: Inline metabolic control in intensified culture Process or cellular state represented: Dynamic nutrient and metabolite state during high-cell-density operation	Measurement or model evidence available: Inline spectroscopy plus associated process measurements Hidden quality state and scale sensitivity: Product-quality consequences may emerge on a different time scale from the controlled metabolite	Qualification test before broader authority: Confirm analytical reliability and demonstrate that control remains valid across the intended culture-density and operating range Permissible consequence and unresolved question: Automatic correction may be appropriate for the qualified metabolic objective; whether the action preserves all relevant CQAs remains a separate question
Manufacturing control case: Soft-sensor-guided manipulation Process or cellular state represented: State variable cannot be measured directly at control frequency	Measurement or model evidence available: Model-derived estimate constructed from available process signals Hidden quality state and scale sensitivity: Estimation error can change with cell line, medium, sensor condition, or operating regime	Qualification test before broader authority: Challenge the estimator under expected perturbations and identify out-of-domain behavior Permissible consequence and unresolved question: Use the estimate only within its demonstrated domain; reduced confidence should contract actuator authority
Manufacturing control case: Adaptive model-predictive control Process or cellular state represented: Future process trajectory under manipulable inputs	Measurement or model evidence available: Learned dynamic model with online parameter updating Hidden quality state and scale sensitivity: Model adaptation may improve process tracking without proving product-quality validity	Qualification test before broader authority: Test constraint handling, model mismatch, disturbance response, and failure behavior independently of nominal optimization Permissible consequence and unresolved question: Permit adaptive control only for objectives and constraints actually qualified; unresolved question is how CQA risk enters the optimization problem
Manufacturing control case: Surrogate-based downstream optimization Process or cellular state represented: Separation state across a multistep unit operation	Measurement or model evidence available: Fast surrogate of a mechanistic or computationally expensive process model Hidden quality state and scale sensitivity: Product purity or recovery objectives are sensitive to the variables and constraints encoded in the surrogate problem	Qualification test before broader authority: Verify surrogate fidelity specifically near decision boundaries and manufacturing constraints Permissible consequence and unresolved question: Computational acceleration is permissible; optimization must not add a quality conclusion absent from the defined objective
Manufacturing control case: Digital-replica-supported integrated control Process or cellular state represented: Linked	Measurement or model evidence available: Multiple unit-operation models combined with process measurements Hidden quality state and scale sensitivity: Error can propagate across interfaces and scale transitions even	Qualification test before broader authority: Qualify component models, interfaces, timing, data lineage, and integrated failure response Permissible consequence and unresolved question: System-level recommendations require system-level evidence; unresolved

upstream and downstream process state	when individual modules perform acceptably	question is when advisory output can acquire autonomous authority
Manufacturing control case: Integrated continuous manufacturing Process or cellular state represented: Persistent material flow across coordinated operations	Measurement or model evidence available: PAT, process models, control logic, and accumulated product evidence Hidden quality state and scale sensitivity: A local control success may have delayed consequences elsewhere in the process train	Qualification test before broader authority: Demonstrate lifecycle control strategy, disturbance propagation, traceability, and product-specific quality protection Permissible consequence and unresolved question: Automation may increase only to the level supported by the integrated qualification package; regulatory acceptability cannot be inferred from technical integration alone

Note: The cross-scale control consequences are proposed synthesis. The table deliberately separates technical control performance from evidence that an action preserves product quality and from evidence supporting broader manufacturing or regulatory transfer.

Distribution shift and model drift

A model can remain mathematically intact while its manufacturing meaning changes. Distribution shift may arise because the cell line, raw-material lot, medium formulation, process scale, equipment configuration, sensor response, sampling procedure, control policy, or product recipe differs from the conditions represented in development data. Online-learning digital-twin frameworks show that models can be updated as dynamic systems evolve, but adaptation itself creates a new inference state whose uncertainty must be assessed rather than assumed away [26].

Sparse production history makes this problem harder. Industrial batch-monitoring methods developed for Low-N environments illustrate that limited historical trajectories can weaken conventional multivariate surveillance and motivate augmentation or alternative learning strategies [27]. For quality-first control, the implication is that model confidence should not depend solely on nominal prediction error. The system should also ask whether the present manufacturing state is adequately represented by the evidence on which that error estimate was obtained. Uncertainty quantification is therefore not an optional display layer added after prediction. Automated biological experimentation increasingly treats uncertainty as part of the data-to-decision pathway because measurement noise, model form, data scarcity, and extrapolation can all influence the reliability of an automated action [28]. In manufacturing control, those uncertainty sources should remain diagnostically separable.

This Perspective consequently proposes four distinct drift states: measurement drift, in which the sensor-to-signal relationship changes; representation drift, in which observed process states depart from the training distribution; relationship drift, in which the mapping from observed state to inferred quality state changes; and decision drift, in which a control policy developed for one operating context is applied where its consequences have not been qualified. These categories are proposed synthesis rather than an established regulatory taxonomy.

The distinctions matter because the appropriate response differs. Sensor recalibration may address measurement drift but not biological relationship drift. Retraining can improve model fit without proving that a modified controller remains safe. Domain extension may justify broader inference but does not automatically justify broader actuator authority. **Table 3** therefore treats drift as a set of manufacturing decision states rather than as a single statistical alarm.

Table 3. Decision states for model drift and reassessment in quality-first bioprocess control

Evidence and limits		Permitted action
Available sensor or model evidence / Unresolved uncertainty	Control decision state / Permissible immediate action	Observation that triggers reassessment / Quality-first synthesis consequence
Available sensor or model evidence: Sensor diagnostics acceptable and current observations resemble qualified training conditions Unresolved uncertainty: Residual model uncertainty within previously examined domain	Control decision state: Stable measurements inside represented operating region Permissible immediate action: Continue qualified automated control	Observation that triggers reassessment: Persistent prediction residual, sensor inconsistency, or approach to domain boundary Quality-first synthesis consequence: Productivity optimization may proceed only within the established quality-compatible envelope
Available sensor or model evidence: Process variables remain plausible	Control decision state: Sensor behavior changes	Observation that triggers reassessment: Orthogonal measurement confirms sensor-

but analytical response or baseline shifts Unresolved uncertainty: Instrument, probe, sampling, or preprocessing change versus true process change	while process trajectory appears familiar Permissible immediate action: Preserve essential control while reducing reliance on the affected channel	specific deviation or genuine process-state change Quality-first synthesis consequence: Measurement validity must be restored before model confidence is interpreted biologically
Available sensor or model evidence: Model produces a prediction but similarity to development data is weak Unresolved uncertainty: Extrapolation error and unknown process-quality relationship	Control decision state: Process enters sparsely represented region Permissible immediate action: Restrict high-consequence optimization and increase independent observation	Observation that triggers reassessment: New representative data establish whether the state belongs to an existing or new operating regime Quality-first synthesis consequence: Prediction availability does not confer control authority outside the qualified domain
Available sensor or model evidence: Updated model fits recent observations better Unresolved uncertainty: Whether improved fit reflects genuine transportability or local overadaptation	Control decision state: Model is recalibrated or retrained after accumulated deviation Permissible immediate action: Use advisory or bounded control until independent challenge succeeds	Observation that triggers reassessment: Holdout, perturbation, or prospective evidence confirms stable performance Quality-first synthesis consequence: Model updating is treated as a change requiring renewed evidence rather than automatic improvement
Available sensor or model evidence: Process model indicates nominal state while direct or higher-proximity quality evidence suggests deviation Unresolved uncertainty: Model failure, delayed biological effect, spatial heterogeneity, or analytical conflict	Control decision state: Process and CQA-relevant evidence become discordant Permissible immediate action: Prioritize product-relevant evidence and narrow automated authority	Observation that triggers reassessment: Root-cause evidence resolves the disagreement Quality-first synthesis consequence: A quality-first architecture permits productivity optimization to yield to unresolved CQA risk
Available sensor or model evidence: Sensor and state models remain acceptable but action-selection logic changes Unresolved uncertainty: New action-response relationship and potential unintended propagation	Control decision state: Control policy itself is modified Permissible immediate action: Limit authority to previously bounded actions	Observation that triggers reassessment: Deliberate challenge shows acceptable disturbance and failure response Quality-first synthesis consequence: Controller qualification cannot be inherited automatically from sensor or prediction-model qualification

Note: The decision states, permissible actions, and reassessment logic are proposed synthesis. They do not constitute numerical drift limits, validated regulatory categories, or universally applicable intervention rules.

Figure 2 places these states inside the control loop so that model updating cannot bypass a central decision boundary separating evidence of improved fit from evidence supporting continued manufacturing authority.




Figure 2. Model drift as a boundary on automated authority in closed-loop bioprocess control

Human oversight

Human oversight is most useful when it changes what the system is permitted to do. A hybrid human-machine model is increasingly considered a practical direction for bioprocess automation because algorithmic speed can be combined with expert judgment when evidence is incomplete, unusual, or consequential [29]. Oversight should therefore be designed as an escalation function rather than a decorative requirement.

The proposed architecture assigns greater human involvement when three conditions increase: uncertainty about the manufacturing state, novelty relative to the qualified domain, and potential consequence for product quality. Routine adjustments inside a well-characterized operating region may require little intervention. In contrast, disagreement between product-proximal evidence and the predictive model, repeated out-of-domain observations, or a proposed controller change should reduce autonomous authority.

This logic avoids two extremes. Requiring manual approval for every adjustment would discard much of the value of closed-loop control. Allowing the model to retain constant authority despite changing evidence would treat automation level as a permanent system property. Quality-first control instead makes autonomy conditional on evidence.

A quality-first AI control architecture

The proposed architecture contains five linked but non-interchangeable functions: measurement, state estimation, CQA-risk interpretation, constrained action selection, and post-action verification. Smart pharmaceutical manufacturing provides the broader technological context for connecting sensors, automation, advanced analytics, and lifecycle quality systems [30]. The distinctive proposal here is the ordering of authority within that connected system.

Productivity is not removed from optimization. It becomes a secondary objective pursued inside a quality-compatible operating envelope. When CQA-relevant evidence is strong and the current process remains within the qualified domain, the system can optimize productivity aggressively within established constraints. When evidence becomes uncertain, productivity remains measurable but its optimization authority contracts.

The architecture also separates three confidence questions. Measurement confidence asks whether the observed signal is analytically reliable. Model confidence asks whether the inference remains supported in the current domain. Decision confidence asks whether the proposed action-response relationship has been qualified at the required consequence level. These dimensions should not be collapsed into one AI confidence value because their failure modes require different corrective actions.

The resulting system is deliberately asymmetric. Evidence can rapidly decrease automation authority when a potentially quality-relevant inconsistency appears, whereas expansion of authority requires additional

qualification. This asymmetry is proposed as a pharmaceutical design principle: uncertainty about productivity may reduce efficiency, but unresolved uncertainty about a CQA-relevant state can change the acceptability of the material being manufactured.

Validation and regulatory qualification

Integrated continuous biomanufacturing frameworks show why qualification should be performed at both component and system levels: connected unit operations share material flow, information, and disturbance propagation, so acceptable local performance does not automatically establish acceptable integrated behavior [31]. The same reasoning applies more strongly when AI models influence control actions.

A quality-first validation strategy should therefore distinguish sensor qualification, prediction-model validation, domain-transport testing, control-policy verification, and integrated failure handling. A model may remain acceptable after sensor recalibration while requiring reassessment after retraining. A predictive model may remain unchanged while a new control policy creates a different risk. A controller may operate correctly in simulation while the measurement latency or disturbance structure of manufacturing invalidates the assumptions under which its performance was demonstrated.

Change control should consequently classify what changed before deciding how much evidence is needed. Recalibration, retraining, domain extension, altered preprocessing, new raw-material ranges, product transfer, scale change, and modification of actuator logic should not be treated as equivalent events. The proposed architecture does not prescribe a universal regulatory pathway; it provides a way to locate the evidentiary burden created by each change.

Prospective qualification is ultimately required to determine whether the quality-first hierarchy improves manufacturing robustness in particular platforms. Relevant evaluation should examine not only nominal prediction accuracy or productivity but also behavior during perturbations, sensor degradation, operating-region transitions, model disagreement, and delayed CQA consequences. Until such testing exists for a specific implementation, the architecture should be regarded as a systems-level design proposition rather than as a validated control standard.

CONCLUSION

AI can make bioprocess control faster, more adaptive, and more information-rich, but computational capability does not determine the correct pharmaceutical objective. A productivity-first controller can optimize what is easy to observe while remaining uncertain about what ultimately defines product acceptability. Quality-first control addresses this mismatch by making CQA-relevant evidence the boundary within which productivity is optimized. The proposed architecture preserves distinctions among measurement, inference, quality interpretation, action, and verification; contracts automation authority under distribution shift or unresolved discordance; and routes recalibration or retraining back through validation rather than directly to control. Its central claim is therefore not that AI should ignore productivity, but that productivity should acquire control priority only after the evidence needed to protect relevant quality attributes is adequate for the intended manufacturing state. Establishing the practical value of this hierarchy will require product-specific, scale-aware, prospectively qualified implementation rather than extrapolation from model performance alone.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Rathore AS, Nikita S, Thakur G, Mishra S. Artificial intelligence and machine learning applications in biopharmaceutical manufacturing. Trends Biotechnol. 2023;41(4):497-510. doi:10.1016/j.tibtech.2022.08.007

2. Khuat TT, Bassett R, Otte E, Grevis-James A, Gabrys B. Applications of machine learning in antibody discovery, process development, manufacturing and formulation: Current trends, challenges, and opportunities. *Comput Chem Eng.* 2024;182:108585. doi:10.1016/j.compchemeng.2024.108585
3. Rathore AS, Sarin D. What should next-generation analytical platforms for biopharmaceutical production look like? *Trends Biotechnol.* 2024;42(3):282-92. doi:10.1016/j.tibtech.2023.08.008
4. Edwards E, Livanos M, Krueger A, Dell A, Haslam SM, Smales CM, et al. Strategies to control therapeutic antibody glycosylation during bioprocessing: Synthesis and separation. *Biotechnol Bioeng.* 2022;119(6):1343-58. doi:10.1002/bit.28066
5. Richter J, Wang Q, Lange F, Thiel P, Yilmaz N, Solle D, et al. Machine learning-powered optimization of a CHO cell cultivation process. *Biotechnol Bioeng.* 2025;122(5):1153-64. doi:10.1002/bit.28943
6. Coolbaugh MJ, Varner CT, Vetter TA, Davenport EK, Bouchard B, Fiadeiro M, et al. Pilot-scale demonstration of an end-to-end integrated and continuous biomanufacturing process. *Biotechnol Bioeng.* 2021;118(9):3287-301. doi:10.1002/bit.27670
7. Rathore AS, Zydney AL, Anupa A, Nikita S, Gangwar N. Enablers of continuous processing of biotherapeutic products. *Trends Biotechnol.* 2022;40(7):804-15. doi:10.1016/j.tibtech.2021.12.003
8. Wu CHR, Chan B, Sarich Z, Duan Y, Chen J, Song JL, et al. Accelerating attribute-focused process and product development through the development and deployment of autonomous process analytical technology platform system. *Biotechnol Bioeng.* 2024;121(4):1256-69. doi:10.1002/bit.28649
9. Khodabandehlou H, Rashedi M, Wang T, Tulsyan A, Schorner G, Garvin C, et al. Cell culture product quality attribute prediction using convolutional neural networks and Raman spectroscopy. *Biotechnol Bioeng.* 2024;121(4):1230-42. doi:10.1002/bit.28646
10. Wan B, Patel M, Zhou G, Olma M, Bieri M, Mueller M, et al. Robust platform for inline Raman monitoring and control of perfusion cell culture. *Biotechnol Bioeng.* 2024;121(5):1688-701. doi:10.1002/bit.28680
11. Thite NG, Yarnell M, Fry TJ, Seefeldt M, Calderon CP, Randolph TW. Unsupervised machine learning-based process analytical tools for near real-time cell morphology analysis during CAR-T cell manufacturing. *Biotechnol Bioeng.* 2025;122(9):2377-88. doi:10.1002/bit.70005
12. Maruthamuthu MK, Rudge SR, Ardekani AM, Ladisch MR, Verma MS. Process analytical technologies and data analytics for the manufacture of monoclonal antibodies. *Trends Biotechnol.* 2020;38(10):1169-86. doi:10.1016/j.tibtech.2020.07.004
13. Kim SJ, Lee M, Kim D, Yoon S, Lee DY. AutoML-driven soft sensors for real-time monitoring of amino acids in mammalian perfusion cultures. *Biotechnol Bioeng.* 2025;122(11):3123-38. doi:10.1002/bit.70036
14. André S, Lagresle S, Da Sliva A, Heimendinger P, Hannas Z, Calvosa É, et al. Developing global regression models for metabolite concentration prediction regardless of cell line. *Biotechnol Bioeng.* 2017;114(11):2550-9. doi:10.1002/bit.26368
15. Rashedi M, Khodabandehlou H, Wang T, Demers M, Tulsyan A, Garvin C, et al. Integration of just-in-time learning with variational autoencoder for cell culture process monitoring based on Raman spectroscopy. *Biotechnol Bioeng.* 2024;121(7):2205-24. doi:10.1002/bit.28713
16. Adebar N, Arnold S, Herrera LM, Emenike VN, Wucherpfennig T, Smiatek J. Physics-informed neural networks for biopharmaceutical cultivation processes: Consideration of varying process parameter settings. *Biotechnol Bioeng.* 2025;122(1):123-36. doi:10.1002/bit.28851
17. Golzarjalal M, Aickelin U, Otte E. Addressing small data challenges in biopharmaceutical development and manufacturing: A mini review of multi-fidelity techniques. *Biotechnol Bioeng.* 2026;123(6):1465-80. doi:10.1002/bit.70213
18. Barón Díaz D, Drommershausen AL, Grünberger A, Holtmann D. Transfer learning approaches in bioprocess engineering: Opportunities and challenges. *Biotechnol Bioeng.* 2026;123(6):1417-31. doi:10.1002/bit.70186
19. Gasparini L, Benedetti A, Marchese G, Gallagher C, Facco P, Barolo M. On the use of machine learning to generate in-silico data for batch process monitoring under small-data scenarios. *Comput Chem Eng.* 2024;180:108469. doi:10.1016/j.compchemeng.2023.108469
20. Sammaknejad N, Lee J, Austria JM, Duenas N, Heiba L, Sridharan G, et al. A scalable deep learning approach for real-time multivariate monitoring of biopharmaceutical processes with no prior product-specific history. *Biotechnol Bioeng.* 2025;122(9):2333-52. doi:10.1002/bit.29039
21. Krausch N, Doff-Sotta M, Cannon M, Neubauer P, Cruz Bournazou MN. Deep learning adaptive model predictive control of fed-batch cultivations. *Comput Chem Eng.* 2025;203:109344. doi:10.1016/j.compchemeng.2025.109344

22. Ding C, Ierapetritou M. Machine learning-based optimization of a multi-step ion exchange chromatography for ternary protein separation. *Comput Chem Eng.* 2024;184:108642. doi:10.1016/j.compchemeng.2024.108642
23. 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
24. Neugebauer P, Zettl M, Moser D, Poms J, Kuchler L, Sacher S. Process analytical technology in downstream-processing of drug substances—A review. *Int J Pharm.* 2024;661:124412. doi:10.1016/j.ijpharm.2024.124412
25. Fisher AC, Kamga MH, Agarabi C, Brorson K, Lee SL, Yoon S. The current scientific and regulatory landscape in advancing integrated continuous biopharmaceutical manufacturing. *Trends Biotechnol.* 2019;37(3):253-67. doi:10.1016/j.tibtech.2018.08.008
26. Rebello CM, Nogueira IBR. Digital twins in chemical engineering: An integrated framework for identification, implementation, online learning, and uncertainty assessment. *Comput Chem Eng.* 2025;200:109178. doi:10.1016/j.compchemeng.2025.109178
27. Tulsyan A, Garvin C, Ündey C. Advances in industrial biopharmaceutical batch process monitoring: Machine-learning methods for small data problems. *Biotechnol Bioeng.* 2018;115(8):1915-24. doi:10.1002/bit.26605
28. 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
29. 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
30. Arden NS, Fisher AC, Tyner KM, Yu LX, Lee SL, Kopcha M. Industry 4.0 for pharmaceutical manufacturing: Preparing for the smart factories of the future. *Int J Pharm.* 2021;602:120554. doi:10.1016/j.ijpharm.2021.120554
31. Coffman J, Brower M, Connell-Crowley L, Deldari S, Farid SS, Horowski B, et al. A common framework for integrated and continuous biomanufacturing. *Biotechnol Bioeng.* 2021;118(4):1735-49. doi:10.1002/bit.27690