Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

DIGITAL-TWIN FRAMEWORK FOR CONTINUOUS MANUFACTURING USING PAT SIGNALS AND CRITICAL QUALITY ATTRIBUTES

Andreas Müller[1]*, Stefan Weber[1], Julia Hoffmann[2], Lukas Schneider[1], Tobias Klein[2]

1. ***Department of Pharmaceutical Data Engineering, Faculty of Pharmacy, Heidelberg University, Heidelberg, Germany.***
2. ***Department of Intelligent Drug Systems, Faculty of Engineering, Technical University of Munich, Munich, Germany.***

ARTICLE INFO

Received:
27 May 2025
Received in revised form:
17 August 2025
Accepted:
17 August 2025
Available online:
28 August 2025

Keywords: Digital twin, Continuous manufacturing, Process analytical Technology, Critical quality attributes, Hybrid modeling, Real-time release

ABSTRACT

Continuous manufacturing is transforming pharmaceutical production by replacing segmented batch operations with integrated, dynamic process trains, which necessitates quality assurance tools that operate continuously rather than relying primarily on end-product release testing. Current control strategies often interpret PAT measurements as isolated trends rather than as part of a connected process state, limiting their ability to anticipate how interacting material attributes, process parameters, and sensor signals shape final product quality. To address this, a digital-twin framework is proposed that ingests multivariate PAT signals, updates a hybrid predictive model, and continuously estimates critical quality attributes in real time, supporting proactive control and real-time release decision-making. The framework integrates a PAT data fusion module, a hybrid physics-informed machine-learning predictor, a digital-twin state estimator, a model drift monitor, and an MES-integrated control advisor, operating as a regulated decision-support architecture rather than an autonomous release mechanism. By enhancing visibility into evolving process quality, enabling early detection of deviations, and helping operators evaluate corrective actions before quality risk becomes material, this approach could also strengthen documentation for real-time release testing through traceable links between sensor data, model predictions, and control recommendations. Overall, an integrated digital-twin framework has the potential to advance the reliability and efficiency of continuous pharmaceutical manufacturing, provided it is implemented with rigorous validation, disciplined model lifecycle management, and alignment with regulatory expectations for predictive quality systems.

To Cite This Article: Müller A, Weber S, Hoffmann J, Schneider L, Tobias Klein T. Digital-Twin Framework for Continuous Manufacturing Using PAT Signals and Critical Quality Attributes. Pharmacophore. 2025;16(4):32-41. https://doi.org/10.51847/sl0aA46GiU

Introduction

The transition from batch manufacturing to continuous pharmaceutical production changes quality assurance from a retrospective release activity into a dynamic process-control challenge. Continuous lines require real-time awareness of material flow, equipment state, and evolving product attributes, because disturbances can propagate across blending, granulation, drying, compression, and coating steps before conventional testing can respond [1]. Reviews of continuous solid-dosage manufacturing emphasize that process models and PAT-enabled monitoring are central to maintaining quality-by-control rather than relying only on terminal inspection [2]. In this setting, a digital-twin framework could provide a structured computational layer that links ongoing process signals to predicted critical quality attributes.

PAT has already shown substantial value for monitoring specific attributes in continuous processes, including blend uniformity through Raman and NIR spectroscopy [3], moisture and particle-size-related changes during granulation and drying [4], and tablet properties during integrated powder blending and compression [5]. However, many implementations still focus on one instrument, one unit operation, or one quality attribute rather than the full multivariate state of a manufacturing line. Continuous manufacturing requires synthesis across spectra, process parameters, residence-time effects, and delayed laboratory

Corresponding Author: Andreas Müller; Department of Pharmaceutical Data Engineering, Faculty of Pharmacy, Heidelberg University, Heidelberg, Germany. E-mail: andreas.mueller@gmail.com

measurements, because product quality emerges from the interaction of these data streams [6]. A framework that treats PAT as a connected state-estimation problem could therefore extend beyond local monitoring toward real-time quality prediction. A digital twin can be understood as a living model that mirrors the current state of a physical process, receives sensor updates, and supports prediction or decision-making under changing operating conditions [7]. Pharmaceutical and biopharmaceutical discussions of digital twins increasingly highlight their potential to integrate process data, mechanistic understanding, and predictive analytics across development and manufacturing contexts [8]. Yet a practical gap remains in architectures that combine physics-based process knowledge with data-driven learning while satisfying regulated expectations for validation, traceability, and lifecycle management [9]. This gap is particularly important for continuous manufacturing, where the model must remain synchronized with the line while producing predictions that are interpretable enough to support quality decisions. This article proposes an AI systems/frameworks perspective on a digital twin that connects real-time PAT data with CQA estimation in continuous pharmaceutical manufacturing. The core thesis is that a hybrid predictive model, combining first-principles process representations with machine-learning residual correction, could improve the robustness of real-time CQA estimates compared with either modeling approach alone [10]. The framework is designed to support closed-loop control recommendations through MES integration while preserving operator oversight and validation discipline [11]. It also positions predicted CQAs and associated uncertainty as components of a real-time release strategy, subject to regulatory evaluation and process-validation expectations [12].

Background

Continuous Manufacturing and Quality-By-Control

Continuous manufacturing replaces sequential batch operations with integrated, dynamically controlled processing, creating opportunities for smaller footprints, improved consistency, and faster response to quality variation. From a quality-by-control perspective, the objective is not simply to test the final product but to maintain the process within a validated design and control space throughout production [1]. Model predictive control has been discussed as a user-oriented strategy for continuous pharmaceutical manufacturing because it can connect process measurements, model predictions, and control actions in a structured way [2]. Regulatory-facing discussions of process models further indicate that model risk, verification, and validation must be considered when predictive models influence manufacturing decisions [12].

Process Analytical Technology in Continuous Lines

PAT instruments such as NIR and Raman spectroscopy provide chemically and physically rich signals that can track blend uniformity, concentration, moisture, and other material attributes during continuous operation. Raman-based monitoring and feedback control have been demonstrated conceptually for continuous powder blending and tableting, showing how spectroscopic signals can be linked to process understanding [3]. NIR-based monitoring has also been used to follow granulation and drying behavior, where spectral information can reflect moisture and particle-related changes relevant to downstream performance [4]. Broader PAT reviews describe these tools as core enablers of continuous process verification when they are embedded in a validated control strategy [13].

Digital Twins in Manufacturing

Digital twins in pharmaceutical manufacturing are typically described as computational representations that combine process data, models, and state-updating mechanisms to represent the current condition of the process [7]. In a manufacturing setting, the twin must not only store data but also maintain synchronization with the physical process and provide predictions that are useful for operational decisions. Recent pharmaceutical perspectives extend the concept across drug development, biomanufacturing, and continuous production, emphasizing the role of real-time data and predictive modeling in transforming process oversight [8]. A process-optimization digital twin further illustrates how computational representations can support scenario analysis and decision-making rather than merely describing historical production data [14].

Hybrid and Physics-Informed Machine Learning Models

Hybrid modeling combines mechanistic process knowledge with data-driven learners, allowing the model to use engineering structure while adapting to empirical deviations. In pharmaceutical continuous manufacturing, this could mean using residence-time distributions, population-balance concepts, or mass-transfer relationships as baseline representations while machine-learning models capture unmodeled nonlinearities [2]. Reviews of artificial neural networks in pharmaceutical PAT note that data-driven models can learn complex relationships from spectra and process variables, but their reliability depends on thoughtful preprocessing, validation, and interpretability [10]. Such hybridization would be expected to improve extrapolation and operator confidence compared with a purely black-box predictor.

Model Maintenance and Real-Time Release

Predictive models used for real-time release must be maintained as manufacturing conditions, raw materials, sensors, and equipment states evolve over time. Regulatory and lifecycle discussions emphasize that models influencing pharmaceutical quality decisions should be verified, validated, monitored, and controlled through formal model-risk frameworks [9]. Practical discussions of PAT and real-time release note that the analytical method, prediction model, and product lifecycle strategy must be considered together rather than as isolated technical components [12]. For a digital twin, this implies that drift detection,

retraining governance, version control, and auditability are core framework requirements rather than optional analytics features.

Framework Architecture Overview

High-Level Digital-Twin Architecture

The proposed architecture consists of four connected layers: a physical layer containing equipment, sensors, and PAT instruments; a data layer that aligns, preprocesses, and fuses heterogeneous signals; a digital model layer that estimates CQAs through hybrid prediction; and a control layer that communicates recommendations to operators and manufacturing systems. This layered view is consistent with pharmaceutical digital-twin concepts that emphasize data integration, model-based prediction, and process-state representation [7]. The control layer should be designed as a decision-support interface, because PAT implementation for advanced control requires practical attention to robustness, validation, and the operational context in which recommendations are acted upon [1]. Such a structure allows new sensors, process steps, or model modules to be added without redesigning the entire twin.

Figure 1 presents the proposed digital-twin architecture linking continuous manufacturing equipment, multivariate PAT streams, validated data fusion, hybrid CQA prediction, drift governance, MES-integrated recommendations, and traceable real-time release support.

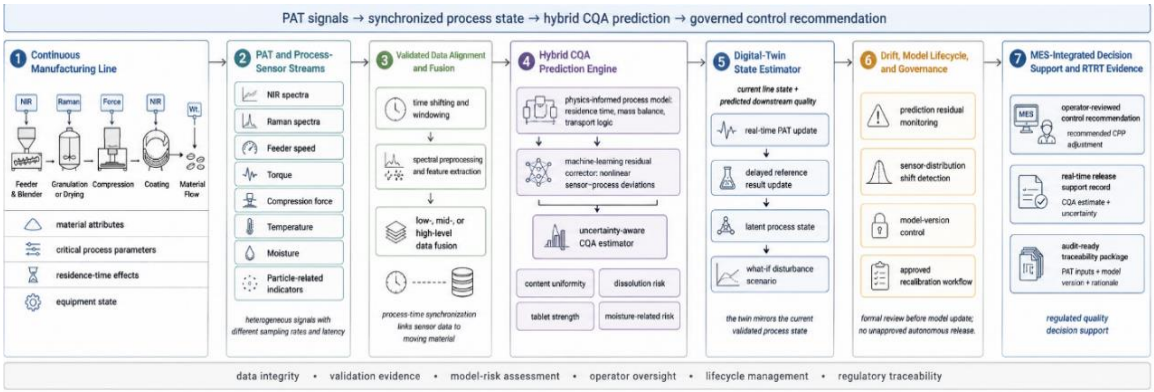

Figure 1. Digital-twin framework for continuous manufacturing using PAT signals and critical quality attributes

Core Input Signals and Outputs

The framework would ingest real-time NIR and Raman spectra, process variables such as feeder speeds and compression force, and derived material indicators such as blend uniformity, moisture content, and particle-size-related features. NIR and Raman spectroscopy have been described as complementary PAT tools for monitoring pharmaceutical solids, with applications across blending, granulation, coating, and tableting [15]. The outputs would include continuous CQA estimates for attributes such as content uniformity, dissolution risk, and tablet mechanical quality, along with uncertainty information and recommended adjustments to critical process parameters. The inclusion of process variables alongside PAT signals reflects the need to connect sensor observations with the operating conditions that generate product quality [6].

Design Principles

The design principles are real-time capability, physics grounding, modularity, data integrity, and continuous recalibration under a validated lifecycle. Real-time capability is needed because continuous production requires the model to update while material is still moving through the line, rather than after a batch is complete [2]. Physics grounding supports interpretability by linking predictions to residence time, material transport, or unit-operation behavior, while modularity allows a twin to extend from blending into granulation, compression, or hot-melt extrusion as process scope changes [14]. Data integrity and traceable recalibration are essential because predictive model outputs may become part of a real-time release or control-support record [12].

Table 1 decomposes the proposed digital twin into functional layers and clarifies how each layer contributes to continuous CQA estimation, control support, validation, and regulated lifecycle governance.

Table 1. Functional Architecture of the Proposed Digital-Twin Framework for Continuous Pharmaceutical Manufacturing

Framework layer	Primary information handled	Core computational function	Quality-decision contribution	Validation or governance requirement
Physical manufacturing layer	Unit operations, material flow, equipment state, feeder behavior, blender conditions, drying conditions, compression force, coating parameters	Establishes the real-world process state that the digital twin must represent	Defines where disturbances originate and how they may propagate downstream into CQA risk	Equipment qualification, sensor placement justification, process-design-space definition, and documented operating limits

PAT acquisition layer	NIR spectra, Raman spectra, moisture indicators, particle-related signals, process sensors, time-stamped equipment parameters	Captures real-time chemical, physical, and process signatures from the moving manufacturing stream	Converts the continuous line from a periodically sampled process into a continuously observed quality system	PAT method qualification, calibration maintenance, sampling-rate documentation, sensor health monitoring, and data-integrity controls
Time alignment and preprocessing layer	Raw spectra, asynchronous process signals, residence-time relationships, delayed reference measurements	Performs time shifting, windowing, interpolation, spectral correction, smoothing, and feature extraction	Ensures that CQA predictions are based on the correct material segment and not on misaligned sensor observations	Preprocessing version control, validated synchronization logic, traceable transformation records, and justification of feature-engineering choices
Data-fusion layer	Preprocessed PAT features, process parameters, material attributes, latent variables, submodel outputs	Integrates heterogeneous sources through low-, mid-, or high-level fusion	Allows CQA estimates to reflect interacting chemical, physical, and operational signals rather than a single isolated trend	Fusion-strategy justification, sensor-compatibility assessment, missing-data handling, and robustness testing across operating conditions
Hybrid predictive-model layer	Physics-based process assumptions, machine-learning residuals, reference CQA measurements, process history	Combines mechanistic prediction with data-driven correction to estimate CQAs and uncertainty	Supports interpretable real-time estimates of content uniformity, dissolution risk, tablet mechanical quality, and moisture-related quality risk	Model-risk classification, training-data representativeness, reference-method linkage, uncertainty calibration, and prospective verification
Digital-twin state-estimation layer	Current sensor state, latent process state, delayed laboratory results, predicted downstream quality	Maintains a synchronized representation of the manufacturing line and updates the process state as new evidence arrives	Enables operators to understand how current upstream behavior may affect downstream CQAs before final testing occurs	State-update traceability, delayed-reference reconciliation, audit trail for prediction revisions, and predefined intervention thresholds
Drift-monitoring and lifecycle layer	Prediction residuals, sensor-distribution changes, operating-region shifts, recalibration evidence, model metadata	Detects model drift, manages recalibration, and controls approved model versions	Prevents silent degradation of prediction reliability during long campaigns, raw-material changes, sensor replacement, or process evolution	Formal change control, model-version documentation, retraining approval workflow, drift-alarm review, and lifecycle performance monitoring
MES-integrated control-advisor layer	CQA estimates, uncertainty intervals, recommended CPP adjustments, operator actions, release-support records	Translates digital-twin outputs into operator-reviewed control recommendations and traceable RTRT evidence	Supports proactive quality decisions while preserving human oversight and avoiding unapproved autonomous release	MES integration validation, role-based access, electronic records compliance, operator rationale capture, and release-decision governance

Pat Data Ingestion and Preprocessing

Multi-Variate Pat Data Acquisition

The PAT ingestion layer would acquire in-line NIR, Raman, and process-sensor streams that may arrive at different sampling rates and with different levels of latency. In continuous tablet manufacturing, spectroscopic monitoring has been used to follow powder blending and tableting behavior, demonstrating the need to align spectral measurements with the movement of material through the line [3]. NIR-based approaches in granulation and drying similarly require synchronization between spectra, equipment conditions, and the evolving state of the product [4]. The framework would therefore align signals on a common process-time basis using validated time-shifting, windowing, or interpolation logic that reflects material residence through each unit operation.

Data Pre-Treatment and Feature Extraction

Spectral data would require preprocessing steps such as scatter correction, baseline handling, smoothing, and feature extraction before being used by the hybrid model. Reviews of NIR applications for blend-uniformity monitoring emphasize that preprocessing and calibration choices strongly influence whether spectral models remain meaningful across process variation [15]. For process sensors, feature extraction could convert raw equipment parameters into fingerprints such as feeder stability, compression-force patterns, torque behavior, or residence-time descriptors. These engineered features would be expected to support more robust CQA prediction because they represent both material signatures and process conditions [13].

Data Fusion – Low-, Mid-, or High-Level

The framework could use low-level fusion by combining preprocessed raw variables, mid-level fusion by concatenating selected latent variables or engineered features, or high-level fusion by combining predictions from separate submodels. PAT data-fusion reviews highlight that different fusion levels are appropriate depending on sensor compatibility, process configuration, and the intended prediction task [6]. Multi-sensor fusion is particularly relevant for powders and granules because no single instrument may fully capture chemical composition, moisture, particle behavior, and equipment state at the same time [16]. A practical digital twin should therefore allow the fusion strategy to be selected and justified during model development rather than fixed as a universal design choice.

Hybrid Predictive Model for CQA Estimation

Physics-Based Process Model

The physics-based component would represent the expected relationship between process settings, material transport, and CQA evolution using simplified mechanistic models. For example, residence-time behavior could connect upstream feeder or blender disturbances to downstream CQA estimates, while population-balance concepts could describe granule-size evolution during wet granulation [2]. In chromatography and other continuous operations, digital-twin work has shown how model-based representations can support process understanding and control design, even when the specific unit operation differs from solid-dosage manufacturing [17]. Within the proposed framework, the mechanistic model would provide a baseline quality estimate and an interpretable structure for what the process should be doing under the current state.

Data-Driven Residual Corrector

The data-driven component would learn the residual between the physics-based estimate and the CQA measured by an accepted reference method, allowing the twin to capture nonlinear effects, sensor interactions, and product-specific deviations. Machine-learning applications in solid oral dosage development show that algorithms can support prediction and decision-making across formulation and process variables when they are carefully validated [18]. Deep-learning approaches have also been explored for continuous solid-dosage manufacturing, suggesting that more flexible models may capture complex relationships in multivariate manufacturing data [19]. In this framework, the residual corrector would be constrained by the mechanistic baseline so that it augments process understanding rather than replacing it with an opaque standalone predictor.

Uncertainty Quantification

The hybrid predictor should return not only a point estimate of each CQA but also an uncertainty range that reflects model confidence under current process and sensor conditions. This is important because real-time release and control decisions require the system to distinguish between a stable prediction and a prediction made under sparse, shifted, or unusual conditions [12]. Model-risk frameworks for pharmaceutical manufacturing indicate that predictive models should be evaluated according to the potential impact of their outputs on product quality and patient risk [9]. In the proposed architecture, uncertainty would guide whether the MES receives a routine control recommendation, an operator-review prompt, or a request for additional reference testing.

Digital Twin Synchronization and Real-Time State Update

Real-Time State Estimation

The digital twin would continuously receive PAT and process-sensor streams and update its internal representation of the manufacturing line as material moves through each unit operation. Blend-uniformity studies in continuous tablet manufacturing show why synchronization between sampled material, spectral measurements, and process location is essential for meaningful state estimation [20]. The hybrid model would operate in a soft-real-time manner, refreshing CQA estimates as new spectra, equipment signals, and derived features become available. This state-update logic would allow the twin to represent not only current sensor values but also the expected downstream quality implications of upstream disturbances.

Handling Time-Delayed Reference Measurements

Reference measurements such as chromatographic assays or offline physical testing would enter the twin with delay, so the framework should incorporate them as corrective evidence rather than as immediate control signals. Continuous manufacturing case studies involving PAT and reference analytics illustrate the need to reconcile in-line predictions with delayed laboratory confirmation during model maintenance and process verification [21]. A Kalman-filter-like update mechanism could adjust the latent process state when delayed results become available, while preserving traceability between the original prediction, the reference result, and any model correction. This approach would help prevent the digital twin from drifting silently away from reference quality measurements over longer manufacturing campaigns.

Process Disturbance Simulation and What-If Scenarios

Beyond real-time monitoring, the digital twin could be used offline to explore process disturbances, material variability, and control strategies before they are implemented on the line. Simulation-oriented digital-twin work in pharmaceutical manufacturing shows how computational process representations can support optimization and risk assessment when linked to process knowledge [14]. A twin could therefore test hypothetical changes in feeder behavior, moisture evolution, blend composition, or compression settings and estimate how those changes would be expected to influence CQAs. Such what-if use

would support process development, deviation investigation, and control-space refinement without implying that simulated outputs replace formal validation evidence.

Model Monitoring, Drift Detection, and Lifecycle Management

Performance Monitoring and Drift Detection

The framework would monitor residuals between predicted CQAs and delayed reference results, while also tracking changes in sensor distributions, preprocessing behavior, and process operating regions. PAT studies in granulation and wet processing demonstrate that changes in water addition, mixing, or material state can alter the relationship between spectral signals and product attributes, making drift monitoring essential [22]. The drift-detection module should therefore evaluate whether prediction errors, latent spectral features, or process fingerprints deviate from the validated operating envelope. When a drift alarm occurs, the system would flag the model for technical review and potential recalibration rather than automatically accepting the updated model into production use.

Automated Model Update and Version Control

When recalibration is justified, the framework could retrain the data-driven residual component using newly approved reference measurements while keeping the physics-based structure and validation rules under formal change control. Machine-learning-enabled NIR workflows emphasize that model selection, sample representativeness, generalizability, and transferability should be treated as managed lifecycle questions rather than one-time development tasks [23, 24]. Each model version should retain metadata describing training sources, preprocessing settings, intended operating range, validation evidence, and approval status. This would create an auditable record that supports regulated review while allowing the digital twin to adapt to evolving products, raw materials, and equipment conditions.

Integration with MES and Closed-Loop Control

Control Recommendations

The MES-integrated control advisor would translate predicted CQA risk into suggested process adjustments, such as modifying feeder behavior, blender conditions, drying intensity, or compression settings within an approved design space. Closed-loop control discussions in continuous tablet manufacturing emphasize that control actions should be grounded in process understanding and should respect the practical constraints of pharmaceutical operations [11]. In-line tablet inspection concepts using machine learning further indicate that AI-based quality signals could support earlier identification of defective or drifting tablet states, provided the operator and validation framework remain central to decision-making [25]. The proposed digital twin would therefore issue recommendations with rationale and uncertainty rather than executing unreviewed autonomous changes.

Supporting Real-Time Release Testing

For real-time release testing, the digital twin would provide a continuous, traceable record of predicted CQAs, PAT inputs, model state, uncertainty, and any operator-approved control recommendations. Practical PAT and RTRT discussions emphasize that predictive release strategies must connect analytical method validation, product lifecycle management, and regulatory expectations [12]. The framework's data-fusion and prediction records could support release decisions by demonstrating that quality was continuously assessed within a validated control strategy rather than inferred only after production. Regulatory acceptance would still require formal evidence that the model, sensors, and control procedures remain suitable for the intended product and manufacturing context.

Table 2 translates digital-twin outputs into manufacturing responses, release-support evidence, and governance safeguards, showing how the framework functions as regulated decision support rather than autonomous quality release.

Table 2. Decision-Support Logic Connecting Digital-Twin Outputs to Manufacturing Actions, Real-Time Release Evidence, and Governance Safeguards

Digital-twin signal or condition	Interpretation within the framework	Recommended manufacturing response	Real-time release or documentation value	Required safeguard before operational use
Stable CQA estimate with low uncertainty	The process appears to remain within the validated operating envelope, and the hybrid model has sufficient confidence in the prediction	Continue routine operation while maintaining standard monitoring frequency	Supports continuous evidence that product quality is being assessed within a validated control strategy	Confirm that the model version, PAT calibration status, and process conditions match the approved operating space
CQA estimate approaching specification boundary	The predicted product attribute may remain acceptable but is moving toward elevated quality risk	Prompt operator review and consider predefined CPP adjustment within approved design space	Creates early-warning documentation before a deviation becomes material	Require operator acknowledgement, rationale capture, and confirmation that recommended action is

				within the validated control strategy
High uncertainty despite acceptable point estimate	The model output is numerically acceptable but confidence is reduced because the current condition is sparse, shifted, or poorly represented	Increase reference sampling, request technical review, or hold reliance on model-only evidence	Prevents false reassurance when a prediction appears acceptable but lacks sufficient support	Define uncertainty thresholds that trigger additional evidence rather than automatic control or release decisions
Disagreement between PAT-based prediction and delayed reference result	The twin's current state estimate may require correction because reference evidence conflicts with prior real-time prediction	Reconcile the latent process state, investigate alignment or calibration issues, and document correction logic	Maintains traceability between original prediction, delayed laboratory result, and any subsequent model update	Preserve audit trail, prohibit silent overwriting of historical predictions, and require quality-system review for repeated discrepancies
Rising prediction residuals across multiple time windows	The residual corrector or physics-informed baseline may no longer represent the process adequately	Initiate drift investigation and evaluate whether recalibration or process intervention is required	Provides documented evidence that model performance is monitored during the product lifecycle	Use predefined drift criteria, distinguish sensor drift from process drift, and prevent automatic deployment of retrained models
Sensor-distribution shift without immediate CQA failure	PAT signals or process fingerprints have moved away from the validated reference distribution even if predicted CQAs remain acceptable	Check sensor status, raw-material attributes, preprocessing behavior, and process conditions	Documents that the control strategy monitors leading indicators rather than only final CQA excursions	Require sensor-health checks, preprocessing audit, and operating-region assessment before continued reliance on predictions
Recommended CPP adjustment within approved design space	The MES advisor has identified a corrective action that is consistent with the validated control strategy	Operator reviews and approves adjustment to feeder speed, blending condition, drying intensity, or compression setting	Demonstrates proactive control based on linked PAT signals, CQA estimates, model rationale, and human oversight	Require role-based approval, electronic record capture, uncertainty display, and rationale for the recommendation
Recommended action outside approved design space	The model suggests that routine control actions may be insufficient or that the process is outside validated conditions	Escalate to deviation management, technical review, or batch/lot segregation procedure	Supports quality-risk documentation by showing that the system did not normalize an unvalidated condition	Block automatic execution, require quality-unit review, and separate decision support from release authorization
Model recalibration candidate identified	Newly approved reference data may improve the residual corrector or update calibration boundaries	Prepare recalibration package but keep current production model unchanged until approved	Creates lifecycle evidence for model maintenance without compromising production traceability	Apply formal change control, validation testing, version comparison, approval workflow, and rollback strategy
Release-support package generated	PAT inputs, model version, predicted CQAs, uncertainty, operator actions, and reference confirmations are assembled into a traceable record	Use as supporting evidence for real-time release evaluation under the approved quality system	Strengthens RTRT documentation by linking sensor evidence, prediction logic, control actions, and governance decisions	Ensure regulatory alignment, data-integrity compliance, method validation, and explicit human release responsibility

Evaluation Strategy

Predictive Accuracy of CQA Estimates

Predictive evaluation should compare digital-twin CQA estimates with accepted reference methods for attributes such as content uniformity, dissolution behavior, and tablet mechanical quality, while avoiding overinterpretation of any single metric. Blend-uniformity monitoring using stream sampling and NIR spectroscopy illustrates how continuous predictions can be compared with reference measurements in a way that reflects process timing and material movement [26]. The evaluation should examine bias, consistency, and robustness across operating conditions conceptually rather than presenting unsupported numerical performance claims. A successful evaluation would show that the model is suitable for decision support within its validated operating space, not that it universally replaces reference testing.

Drift Detection Effectiveness

Drift-detection evaluation should assess whether the system can identify meaningful changes in sensors, materials, or process behavior before those changes compromise quality decisions. Data-fusion work using material databases and process-model development for high-shear wet granulation suggests that broader information sources can improve the ability to interpret

changes in process behavior [27]. The drift monitor should therefore be evaluated using planned stress scenarios, historical deviations, and controlled recalibration cases, with attention to false reassurance as well as excessive alarm generation. Its purpose would be to maintain trust in the digital twin over time by ensuring that the model's operating context remains aligned with its validated assumptions.

Control Improvement and Efficiency

Control evaluation should focus on whether the digital twin would support more timely, interpretable, and quality-relevant decisions than conventional monitoring alone. Studies combining NIR, Raman, and data-fusion concepts for tablet-quality prediction show how multi-source information can strengthen interpretation of complex process states [28]. The assessment should consider process capability conceptually, deviation handling, operator workload, sampling strategy, and documentation quality, without claiming numerical improvements that were not experimentally demonstrated. In an AIF article, the goal is to define how such improvements should be evaluated in future implementation studies rather than to report fabricated performance outcomes.

Limitations

Need For High-Quality Historical Data to Build the Twin

The proposed framework depends on representative historical and ongoing data that cover raw-material variability, process disturbances, equipment states, and product-specific behavior. Machine-learning applications in solid oral dosage development indicate that predictive usefulness depends strongly on the quality and relevance of the data used to train and validate the model [18]. New products, early development programs, or startup phases may lack sufficient process diversity to support a reliable residual corrector. In those cases, the twin should begin as a constrained decision-support tool and expand only as validated evidence accumulates.

Regulatory Hurdles for Real-Time Release

Regulatory acceptance of AI-augmented predictive quality systems remains an evolving area, especially when model outputs influence release decisions or process interventions. Model-risk and validation frameworks for pharmaceutical manufacturing emphasize the need to define model purpose, assess impact, document assumptions, and maintain lifecycle controls [9]. Even when a digital twin is technically plausible, its use in RTRT would require alignment among PAT validation, process validation, data integrity, change management, and operator governance. The framework should therefore be viewed as a pathway toward regulated real-time release rather than as an automatic substitute for established quality systems.

Conclusion

The proposed digital-twin framework defines a conceptual architecture for continuous pharmaceutical manufacturing that connects PAT signals, process parameters, hybrid predictive modeling, and CQA estimation. It treats quality as an evolving process state rather than as a property discovered only through end-product testing. By maintaining a synchronized representation of the manufacturing line, the framework could help operators understand how current process behavior may affect downstream product quality.

Its main strength is the integration of real-time quality visibility, physics-informed AI, model lifecycle management, and MES-linked decision support. The physics-based component would provide interpretability, while the data-driven residual corrector would capture product- and process-specific deviations. The drift monitor and version-control layer would help keep the model aligned with regulated manufacturing expectations. Together, these features create a practical pathway toward predictive control and real-time release support.

Important challenges remain before such a framework can be adopted broadly. The twin would require high-quality historical data, representative reference measurements, reliable sensor maintenance, and disciplined validation across the product lifecycle. Regulatory maturity will also be necessary so that AI-supported quality predictions can be assessed consistently and transparently. Cybersecurity and data-integrity protections would be essential because the framework depends on connected instruments, manufacturing systems, and decision-support interfaces.

Future progress should come from collaborative demonstration projects involving manufacturers, technology providers, academic groups, and regulatory agencies. These projects should establish validation precedents for hybrid models, uncertainty-aware CQA prediction, drift management, and MES-integrated control recommendations. They should also clarify how digital-twin evidence can support real-time release while preserving human oversight and formal quality governance. Such collaboration could make digital-twin-enabled continuous manufacturing a credible and auditable foundation for future pharmaceutical production.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Sacher S, Poms J, Rehrl J, Khinast JG. PAT implementation for advanced process control in solid dosage manufacturing–a practical guide. Int J Pharm. 2022;613:121408.
2. Jelsch M, Roggo Y, Kleinebudde P, Krumme M. Model predictive control in pharmaceutical continuous manufacturing: A review from a user's perspective. Eur J Pharm Biopharm. 2021;159:137-42.
3. Nagy B, Farkas A, Gyürkés M, Komaromy-Hiller S, Démuth B, Szabó B, et al. In-line Raman spectroscopic monitoring and feedback control of a continuous twin-screw pharmaceutical powder blending and tableting process. Int J Pharm. 2017;530(1-2):21-9.
4. Pauli V, Roggo Y, Kleinebudde P, Krumme M. Real-time monitoring of particle size distribution in a continuous granulation and drying process by near infrared spectroscopy. Eur J Pharm Biopharm. 2019;141:90-9.
5. Nagy B, Farkas A, Magyar K, Démuth B, Nagy ZK, Marosi G. Spectroscopic characterization of tablet properties in a continuous powder blending and tableting process. Eur J Pharm Sci. 2018;123:10-9.
6. Casian T, Nagy B, Kovács B, Galata DL, Hirsch E, Farkas A. Challenges and opportunities of implementing data fusion in process analytical technology—A review. Molecules. 2022;27(15):4846.
7. Chen Y, Yang O, Sampat C, Bhalode P, Ramachandran R, Ierapetritou M. Digital twins in pharmaceutical and biopharmaceutical manufacturing: a literature review. Processes. 2020;8(9):1088.
8. Maharjan R, Kim NA, Kim KH, Jeong SH. Transformative roles of digital twins from drug discovery to continuous manufacturing: Pharmaceutical and biopharmaceutical perspectives. Int J Pharm X. 2025:100409.
9. Bideault G, Scaccia A, Zahel T, Landertinger RW, Daluwatte C. Verification and validation of computational models used in biopharmaceutical manufacturing: Potential application of the ASME verification and validation 40 standard and FDA proposed AI/ML model life cycle management framework. J Pharm Sci. 2021;110(4):1540-4.
10. Nagy B, Galata DL, Farkas A, Nagy ZK. Application of artificial neural networks in the process analytical technology of pharmaceutical manufacturing—a review. AAPS J. 2022;24(4):74.
11. Malevez D, Copot D. From batch to continuous tablet manufacturing: a control perspective. IFAC Pap OnLine. 2021;54(15):562-7.
12. Panzitta M, Calamassi N, Sabatini C, Grassi M, Spagnoli C, Vizzini V, et al. Spectrophotometry and pharmaceutical PAT/RTRT: Practical challenges and regulatory landscape from development to product lifecycle. Int J Pharm. 2021;601:120551.
13. Kim EJ, Kim JH, Kim MS, Jeong SH, Choi DH. Process analytical technology tools for monitoring pharmaceutical unit operations: a control strategy for continuous process verification. Pharmaceutics. 2021;13(6):919.
14. Gerogiorgis DI, Castro-Rodriguez D. A digital twin for process optimisation in pharmaceutical manufacturing. In: Comput Aided Chem Eng. 2021;50:253-8.
15. Velez NL, Drennen JK, Anderson CA. Challenges, opportunities and recent advances in near infrared spectroscopy applications for monitoring blend uniformity in the continuous manufacturing of solid oral dosage forms. Int J Pharm. 2022;615:121462.
16. Asachi M, Camargo-Valero MA. Multi-sensors data fusion for monitoring of powdered and granule products: Current status and future perspectives. Adv Powder Technol. 2023;34(7):104055.
17. Tiwari A, Masampally VS, Agarwal A, Rathore AS. Digital twin of a continuous chromatography process for mAb purification: Design and model-based control. Biotechnol Bioeng. 2023;120(3):748-66.
18. Lou H, Lian B, Hageman MJ. Applications of machine learning in solid oral dosage form development. J Pharm Sci. 2021;110(9):3150-65.
19. Roggo Y, Jelsch M, Heger P, Ensslin S, Krumme M. Deep learning for continuous manufacturing of pharmaceutical solid dosage form. Eur J Pharm Biopharm. 2020;153:95-105.
20. Sierra-Vega NO, Román-Ospino A, Scicolone J, Muzzio FJ, Romañach RJ, Méndez R. Assessment of blend uniformity in a continuous manufacturing process. Int J Pharm. 2019;560:322-33.
21. Roggo Y, Pauli V, Jelsch M, Pellegatti L, Elbaz F, Ensslin S, et al. Continuous manufacturing process monitoring of pharmaceutical solid dosage form: A case study. J Pharm Biomed Anal. 2020;179:112971.
22. Koyanagi K, Ueno A, Sasaki T, Otsuka M. Real-time monitoring of critical quality attributes during high-shear wet granulation process by near-infrared spectroscopy: effect of water addition and stirring speed on pharmaceutical properties of the granules. Pharmaceuticals. 2022;15(7):822.
23. Ali H, Muthudoss P, Ramalingam M, Kanakaraj L, Paudel A, Ramasamy G. Machine learning–enabled NIR spectroscopy. Part 2: workflow for selecting a subset of samples from publicly accessible data. AAPS PharmSciTech. 2023;24(1):34.
24. Ali H, Muthudoss P, Chauhan C, Kaliappan I, Kumar D, Paudel A, et al. Machine learning-enabled NIR spectroscopy. Part 3: hyperparameter by design (HyD) based ANN-MLP optimization, model generalizability, and model transferability. AAPS PharmSciTech. 2023;24(8):254.

25. Kim SH, Han SH. Development of an intelligent tablet press machine for the in-line detection of defective tablets using machine learning and deep learning models. Pharmaceutics. 2025;17(4):406.
26. Rangel-Gil RS, Nasrala-Álvarez JM, Romañach RJ, Méndez R. Blend uniformity monitoring in a continuous manufacturing mixing process for a low-dosage formulation using a stream sampler and near infrared spectroscopy. Int J Pharm. 2024;661:124478.
27. Wang Z, Cao J, Li W, Wang Y, Luo G, Qiao Y, et al. Using a material database and data fusion method to accelerate the process model development of high shear wet granulation. Sci Rep. 2021;11(1):16514.
28. Szabó-Szőcs B, Ficzere M, Peterfi O, Galata DL. Simultaneous prediction of the API concentration and mass gain of film coated tablets using near-infrared and Raman spectroscopy and data fusion. Int J Pharm. 2025;668:124957..