Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Governed Synchronization of Adaptive Digital Twins for Industrial Biomanufacturing: Application to Penicillin Fermentation

Suchandan Banerjee[1]
National Institute of Technology Durgapur
Shrestha Giri
National Institute of Technology Durgapur
Mithu Paul
National Institute of Technology Durgapur
Sandip Kumar Lahiri
sklahiri.che@nitdgp.ac.in
National Institute of Technology Durgapur

Research Article

Keywords: Adaptive Digital Twin, Biomanufacturing, Multi-Gene Genetic Programming (MGGP), Penicillin Fermentation, Governed Synchronization

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11603582/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Governed Synchronization of Adaptive Digital Twins for Industrial Biomanufacturing: Application to Penicillin Fermentation

Suchandan Banerjee[1], Shrestha Giri[2], Mithu Paul[2], Sandip Kumar Lahiri*

Department of Chemical Engineering, National Institute of Technology Durgapur, Durgapur, West Bengal 713209, India

Corresponding author: sklahiri.che@nitdgp.ac.in*

Abstract

Digital twins are increasingly used in biomanufacturing to support real-time monitoring and decision-making. However, maintaining long-term synchronization between the digital twin and the physical twin remains a significant challenge because biological processes continuously evolve due to batch-to-batch variability, process drift, and measurement uncertainty. Existing approaches rely on repeated model retraining or high-dimensional adaptive models, increasing computational complexity and reducing interpretability and industrial deployability. This work proposes a Governed Adaptive Digital Twin (GADT) framework for industrial penicillin fermentation that relies on a validated symbolic process model synchronized through a low-dimensional governed adaptation layer. The structural model is developed offline using multi-gene genetic programming (MGGP), while online synchronization is achieved by adaptively updating only two interpretable parameters, gain and bias, using normalized prediction error. The adaptation process incorporates scheduling, innovation filtering, guardrail constraints, parameter bounding, and batch lifecycle management to ensure robust industrial operation. The proposed framework was evaluated using pseudo-online streaming data generated from 90 simulated fed-batch fermentation campaigns based on the industrial-scale IndPenSim benchmark, labelled as representing approximately three years of simulated production campaigns in the IndPenSim benchmark. Compared with a conventional digital twin, the proposed GADT achieved higher prediction accuracy and trajectory-level synchronization while effectively compensating for long-term process drift without recursive model correction or structural model retraining. Ablation studies further demonstrate that the proposed governance mechanisms are essential for stable and physically consistent adaptation. The results demonstrate that governed synchronization, rather than continuous structural model evolution, enables sustained digital twin performance and provides an interpretable, computationally, industrially deployable architecture for adaptive digital twins in biomanufacturing.

Keywords Adaptive Digital Twin, Biomanufacturing, Multi-Gene Genetic Programming (MGGP), Penicillin Fermentation, Governed Synchronization

1. Introduction

Because their behaviour continuously evolves during operation, industrial biomanufacturing processes are difficult to mimic with static predictive models [1]. The observed challenge of changing conditions is especially pronounced in fed-batch fermentation, consisting of changing levels of microbial activity, substrate utilisation, precursor consumption and product formation during different physiological phases [2, 3]. Regulation of product formation in penicillin (a secondary metabolite) fermentation is strongly influenced by the changing physiological state of the organism. When nominal operating recipes are repeated, batch-to-batch variation can occur because of the variation in the quality of the inoculum, the condition of raw materials, oxygen transfer, availability of precursors, and environmental disturbances. Due to this, dependable long-horizon prediction of penicillin concentration is still a challenging problem for process monitoring and decision support in industry. This creates another important challenge because direct online measurements of product concentration are scarce [4]. Although operating variables such as aeration rate, broth volume, feed rates, and base addition can be monitored continuously, penicillin concentration is generally determined through laboratory analysis or periodic pseudo-online quality measurements. Consequently, the digital twin must operate using sparse and potentially delayed quality information rather than continuous product-concentration measurements. Due to this measurement limitation, synchronizing the physical process and its digital counterpart is particularly difficult in long fermentation batches[1, 4].

Conventional digital twins and soft sensors are often developed offline from historical process data and then deployed as fixed, non-adaptive predictive models [5]. Although such models can predict time-varying process trajectories, their model structure and parameters are not updated online unless an adaptation mechanism is explicitly introduced. Although these models may work well at nominal conditions, they implicitly assume that the relationship between process variables and product concentration will remain sufficiently stationary in operation. This assumption is often disregarded in industrial fermentation. Prediction trajectories may deviate from measured process behaviour because of batch-to-batch variability, measurement bias, unmodelled disturbances, and possible process changes. The resulting mismatch can reduce the reliability of the model for monitoring and decision support. In the present study, these discrepancies are evaluated through pseudo-online replay and are not, by themselves, taken as proof of chronological process drift[6]. To allow for a solution towards this end, adaptive modelling strategies are being suggested to update model parameters, retrain data-driven models, or institutionalize online learning algorithms as new measurements are available[2]. Although these approaches enhance accuracy, they also come with practical drawbacks. High-dimensional online adaptation can be costly, difficult to interpret, and hard to verify. Furthermore, reiterative adjustment of the validated model structure may deteriorate lifecycle traceability, which is an important demand in industrial and regulated biomanufacturing contexts[1]. This therefore calls for adaptive digital twin approaches that can sustain synchronization without the constant need to redevelop or retrain the basic structural model.

The Governed Adaptive Digital Twin Framework (GADT) for an industrial penicillin fermentation is proposed. The main idea is to separate structural process modelling from on-line synchronization. Initially, an understandable symbolic process model is constructed offline by multigene genetic programming (MGGP)[7, 8]. This model serves as the invariant structural representation of the fermentation process and remains unchanged during deployment. Following this, an online synchronization procedure is executed through a corrective layer with a low dimensionality that updates only 2 parameters. Specifically, the gain and the bias are the two physically interpretable parameters that undergo this update[3]. The gain offsets multiplicative changes in process productivity and trajectory magnitude. The bias offsets any additive change

caused by measurement bias, unmodelled disturbances, or accumulation of residual error. To ensure stable industrial operation, the adaptive layer is explicitly governed. The proposed governance mechanisms include scheduled adaptation, innovation filtering, error-triggered responsiveness, parameter bounds and batch-wise reinitialization. These mechanisms restrict adaptation to physically meaningful and operationally acceptable behaviour, thereby preventing unstable parameter evolution and reducing sensitivity to abnormal measurements. (Fig. 1) summarizes the motivation for the proposed approach by contrasting drift-prone static digital twins with the governed synchronization strategy adopted in this work. The suggested framework was evaluated through the use of pseudo-online streaming data with respect to 90 simulated fed-batch fermentation campaigns from the industrial-scale IndPenSim benchmark over nearly three years of routine production. The evaluation setup in this paper is more representative of practical deployment than the usual offline testing because the quality values are introduced one after another during batch operation. Hence, this study evaluates not only prediction quality, but also the capacity of the digital twin to achieve trajectory-level synchronization over extended fermentation horizons.

This study makes four main contributions to adaptive digital twins for industrial biomanufacturing in four essential ways. To begin, a symbolic MGGP model that is interpretable is created as a fixed structural digital twin[8]. Next, a regulated alignment-bias synchronization strategy, which does not require a structure retrain, is proposed to govern the structural model with the evolving physical process. Third, a pseudo-online evaluation framework is established through the use of multi-year industrial fermentation data. In conclusion, the results reveal that compliance-driven synchronization, rather than constant model evolution, enables sustained digital twin performance. Moreover, a realistic architecture for adaptive digital twins for industrial biomanufacturing is presented.

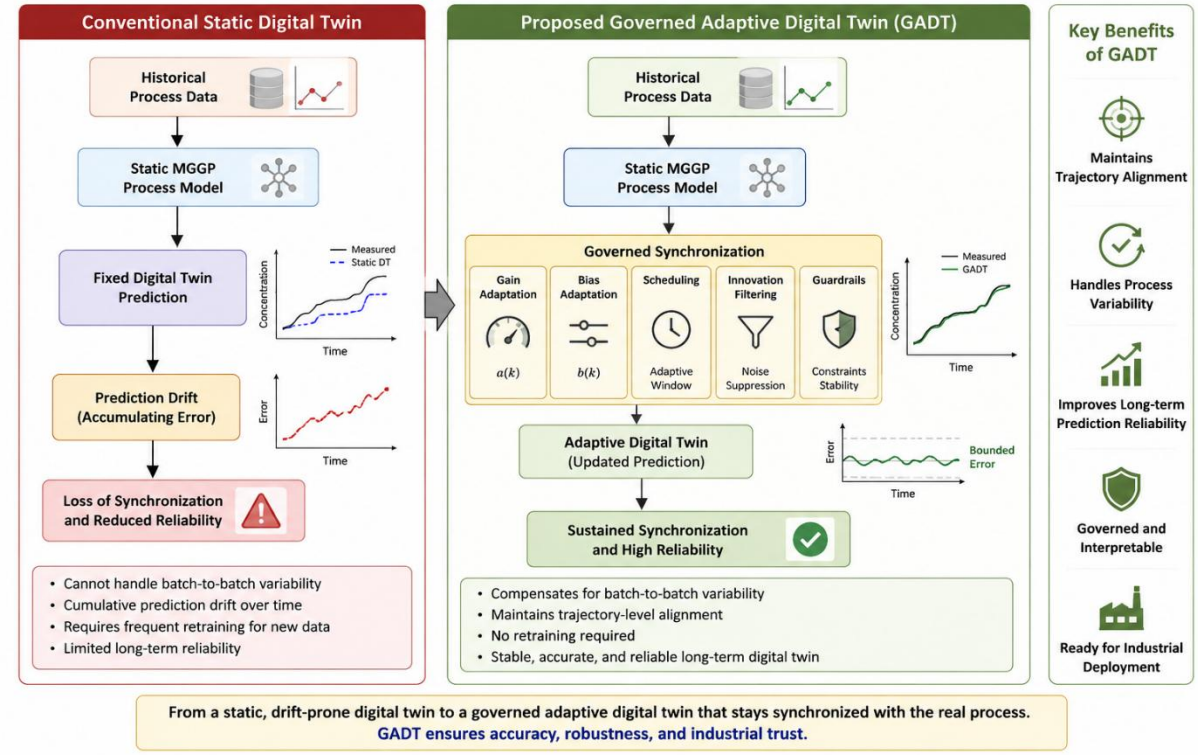

Fig. 1 Evolution from conventional static digital twins to the proposed Governed Adaptive Digital Twin (GADT). The framework preserves the validated structural model while maintaining long-term synchronization through governed online adaptation

2. Process Description and Data Context

2.1 Industrial Penicillin Fermentation Process

Industrial penicillin fermentation was selected because it represents a complex biomanufacturing process characterized by long time horizons, nonlinear dynamics, sparse product measurements, and substantial batch-to-batch variability. The process is managed in fed-batch mode involving continuous evolution of microbial growth, utilization of substrate, feeding of precursor and formation of product with time. The synthesis of penicillin is mainly a secondary metabolism, depending more on the physiological state of a microorganism than on the operating conditions imposed on it[9].

Although the same production recipes are sometimes followed, differences in inoculum quality, raw material properties, oxygen transfer, precursor availability, and environmental disturbances can give rise to slightly different batch trajectories. Due to these variations, a fixed, non-adaptive model cannot easily predict penicillin concentration over time. Moreover, the batches of the fermenting can extend over days, allowing a small prediction error to accumulate to larger trajectory-level errors. As such, the procedure lends itself for assessing whether an adaptive digital twin achieves and maintains synchronization with a real-life industrial process [10].

2.2 Industrial Dataset and Process Variables

The digital twin was developed and evaluated using the publicly available IndPenSim industrial-scale penicillin fermentation benchmark dataset[10]. The dataset comprises 90 simulated fed-batch fermentation campaigns labelled as representing approximately three years of

simulated production campaigns in the IndPenSim benchmark, with routinely measured process variables sampled at 0.2 h intervals. Among the dataset are 90 fed-batch fermentation runs that last about 230 h for each batch. The sampling of routinely measured process variables took place every 0.2 hours to cover the complete batch[10].

The 90 runs are treated in this study as separate simulated fermentation campaigns. The dataset's campaign or year labels are not interpreted here as independent evidence of a chronological process-drift mechanism. Accordingly, the analyses below focus on variation across the evaluated campaigns and on the performance of the proposed synchronization method under the specified pseudo-online replay protocol. The soft sensor to predict penicillin concentration was formulated as a digital twin using routinely available process measurements. Twelve routinely measured process variables were initially considered as candidate inputs for symbolic model development: time, sugar feed rate, phenylacetic acid (PAA) feed rate, aeration rate, dumped broth flow, dissolved oxygen, vessel volume, pH, temperature, CO_2 off-gas, O_2 off-gas, and substrate concentration. The final input set was determined through the repeated MGGP variable-selection procedure described in Section 3.4; therefore, the candidate input set should not be confused with the variables retained in the final symbolic expression. Penicillin concentration expressed as g L^{-1} is the model output. Observed values during the production campaign range from about 0 to 35 g L^{-1}. As a result, the dataset captures temporal evolution within the batch and variability between batches at the industry level[9]. The main features are described in Table 1.

Table 1 Summary of the industrial penicillin fermentation dataset

Dataset characteristic	Description
Dataset source	IndPenSim industrial-scale penicillin fermentation benchmark
Number of fermentation batches	90
Sampling interval	0.2 h
Approximate batch duration	230 h
Observations per batch	Approximately 1150
Candidate input variables	12
Output variable	Penicillin concentration
Total variables per observation	13
Approximate total observations	90 × 1150 = 103,500

2.3 Measurement Constraints and Digital Twin Requirements

One of the main practical challenges in penicillin fermentation is that product concentration is usually not available online continuously. Process variables such as flow rates, vessel volume, and aeration can be recorded continuously; however, penicillin concentration is generally determined using laboratory methods or pseudo-online quality measurements. Consequently, the digital twin needs to operate on sequential and intermittent quality information rather than constant product feedback. However, in this study, penicillin observations were presented to the synchronization layer at the same 0.2 h interval as the process measurement [9].

In this work, the measured penicillin concentrations are treated as pseudo-online observations that become available sequentially during batch operation. The assumption made enables the assessment of the adaptive digital twin under conditions that are representative of the real-world application, where synchronization updates can only be made when new quality information becomes available [11]. For this evaluation, penicillin concentration observations were supplied sequentially at the stated 0.2 h interval. This is a simulation assumption of the present pseudo-online protocol, rather than evidence that laboratory penicillin measurements are routinely available at that frequency in industrial practice. The consequences of delayed, irregular, or missing quality measurements therefore remain to be assessed under more realistic online deployment conditions.

The principal design requirements of the proposed framework are these traits. The digital twin should be accurate at long-horizon predictions, robust to batch-to-batch variations, sparse quality measurements, and computationally efficient and interpretable. Since a purely fixed, non-adaptive model cannot rectify the accumulated process drift, it will not fully satisfy these requirements. On the contrary, continuously retrained models could alter traceability and validation process. Consequently, a synchronization layer with governing rules corrects discrepancies specific to batches that may or may not occur during operations[12].

3. Development of the Structural Digital Twin Model

3.1 Design Philosophy for the Structural Digital Twin

The proposed adaptive digital twin features an invariant knowledge layer called the structural model. The aim for this case is to characterise the most relevant relationship between the fermentation variables that are measured regularly and penicillin concentration, while being stable, interpretable and computationally cheap for long-term industrial use. In the proposed framework, the structural model is being devised offline and not altered during online operation. The design retains historical data process knowledge and avoids the validation challenge of constantly evolving model structure[13].

These are particularly important for industrial biomanufacturing, which suffers from prediction drift in deployment through long batch durations, biological variability, and limited quality measurements. Using a very complicated black-box model may augment the offline fitting accuracy but it increases opacity and obfuscates the management of the lifecycle complications. In contrast, a very simplistic model would be unable to show the nonlinear behaviour of product formation. Thus, the engineering interpretability must be balanced with the predicting capability[14].

Multigene genetic programming (MGGP) was selected for this purpose because it produces explicit symbolic equations while maintaining nonlinear modeling capability. The final model can be checked, evaluated quickly, and utilized as a fixed representation of the fermentation process. This illustrative demonstration also separates online adaptation from structural modelling[8].

Thus, the MGGP model created in this study is the structural underlay of the governed adaptive digital twin. Synchronization errors are not resolved by retraining or altering this symbolic model at deployment. They are instead compensated via a separate low-dimensional governed synchronization layer, which will be introduced in the upcoming sections. The aforementioned separation between variants structural modelling and online adaptive synchronisation is the main design principle of the framework[15].

3.2 Symbolic Model Formulation

The objective of the structural model is to estimate the penicillin concentration from routinely measured process variables that are continuously available during industrial fermentation. The model inputs are defined in Table 2

Table 2 The model input variables of fermentation dataset

Variables	Process Variables	Unit
Input variables		
x_1	**Time**	**h**
x_2	**Sugar feed rate**	**L/h**
x_3	**Phenyl acetic acid (PAA) feed rate**	**L/h**
x_4	**Aeration rate**	**L/h**
x_5	**Dumped broth flow**	**L/h**
x_6	**Dissolve Oxygen,**	**mg/L**
x_7	**Vessel volume**	**L**
x_8	**pH**	**(-)**
x_9	**Temperature**	**K**
x_{10}	**CO2 off gas**	**%**
x_{11}	**O2 off gas**	**%**
x_{12}	**Substrate Concentration**	**g/L**
Output variable		
y	**Penicillin concentration**	**g/L**

These variables were selected based on process knowledge and data availability, and collectively represent the key manipulated and state-dependent factors governing oxygen transfer, carbon availability, pH regulation, and precursor supply. The model output is the penicillin concentration, expressed in g/L, with observed values ranging from 0 to approximately 35 g/L across the dataset.

Input variables as per table 2

$$\mathbf{x}(t) = \boldsymbol{x_1}, \boldsymbol{x_2}, \boldsymbol{x_3}, \ldots, \boldsymbol{x_{12}}$$

The model output,

$$y(t),$$

represents the penicillin concentration ($\mathrm{g\,L^{-1}}$).

Within the MGGP framework, the predicted penicillin concentration is expressed as

$$\hat{y}_{\mathrm{MGGP}}(t) = \sum_{i=1}^{M} c_i \, \phi_i\big(\mathbf{x}(t)\big)$$

where $\phi_i(\cdot)$ denote nonlinear symbolic basis functions automatically evolved through genetic programming, c_i are linear weighting coefficients, and M is the number of genes comprising the final symbolic model.

The model incorporates nonlinear representation ability in a simple analytical expression. The complex interactions of fermentation variables will be captured through the nonlinear basis functions and the linear coefficient estimation allows numerical stability and efficient model evaluation during online operation[16].

3.3 Offline Development of the Structural Model

The MGGP model was created entirely offline using the industrial dataset described in Section 2. Model development was organized at the fermentation-batch level to prevent observations from the same batch from appearing in more than one data subset, thereby minimizing information leakage between batches. The 90 fermentation batches (Batches 1–90) were divided into training, validation, and test subsets using a 70:15:15 allocation. Accordingly, 63 batches (70%) were assigned to training, while the remaining 27 batches were divided between validation and testing, comprising 13 and 14 batches, respectively. The training subset was used for fitting the MGGP model, the validation subset for model selection and assessment during development, and the independent test subset for evaluating the predictive performance of the selected model [10].

To ensure reproducibility, the MGGP search was configured to use a population of 500 individuals and run for a maximum of 350 s. The configuration did not set any maximum limit for generations or iterations. The termination based on fitness was set to a value of 0.0001. The set of candidate-expression primitives included addition, subtraction, multiplication, squaring, exponentiation, division, logarithmic transformation, square root, cubic transformation, negation, absolute value, negative exponential, and one three-argument multiplication operator and one three-argument addition operator. The genetic search employed tournament selection with a tournament size of 25, a Pareto-selection parameter of 0.7, and an elite fraction of 0.7. The constant-initialization parameter was set to 0.5, and the maximum number of genes per individual was restricted to six. The configuration specified one independent run. Crossover and mutation probabilities, the precise fitness-function definition, the symbolic-complexity metric, and the final Pareto-front candidate-selection rule require confirmation from the corresponding implementation files before they can be reported.

3.4 Robust Structural Variable Identification

The original data set obtained from industry had more routinely measured process variables compared with those that were finally the requisite for the structural digital twin. Having all the variables can increase the symbolic complexity and the lack of interpretability. Thus, a robustness-oriented variable identification strategy was utilized. The MGGP Algorithm was run independently 100 times, with a different randomly initialized population each time, but same training data and evolutionary settings. After each run, the variables in the final symbolic expression were captured. The frequency range of a variable's occurrence across independent MGGP runs was taken as an indicator of structural relevance of the variable. A variable selected with regularity is a robust contributor to prediction of penicillin, while a variable selected only rarely is weak or context-dependent. Table 3 summarizes the occurrence frequencies of all findings. CO_2 Off-gas (x_{10}), Dissolved Oxygen(x_{6}), O_2 Off-gas (x_{11}), Phenyl acetic Acid (PAA) Flow Rate (x_3), Sugar Feed Rate (x_2), Substrate Concentration (x_{12}), pH(x_8), Aeration Rate (x_4) occurred most often in a symbolic model showing their dominance. Although vessel volume was included among the candidate process variables, it was not selected in any of the 100 independent MGGP runs (0% occurrence) and was therefore excluded from the final structural input set. Variables with low selection frequency were excluded. The repeated-discovery methodology enhances confidence in a presented set of variables since it supports consistency of structure rather than the variance of a single evolutionary run.

Table 3 Robust relevance of candidate input variables identified from repeated MGGP model

Candidate input variable	Frequency of appearance in equations	Percentage (%)
CO_2 Off-gas (x_{10})	97	97
Dissolved Oxygen(x_{6})	93	93
O_2 Off-gas (x_{11})	93	93
Phenyl acetic Acid (PAA) Flow Rate (x_3)	81	81
Sugar Feed Rate (x_2)	67	67
Substrate Concentration (x_{12})	27	27
pH(x_8)	9	9
Aeration Rate (x_4)	8	8
Vessel Volume (x_7)	0	0.00
Dumped Broth Flow (x_5)	0	0.00
Temperature (x_9)	0	0.00
Time (x_1)	0	0.00

The zero selection frequency of dumped broth flow indicates that this variable was not retained in the final symbolic expressions generated by the 100 MGGP runs under the candidate-variable set, primitive set, training data, and selection procedure used in this study. This result should be interpreted as a model-selection outcome rather than evidence that dumped broth flow has no physical influence on penicillin fermentation. Its contribution may be indirect, redundant with other measured variables, or insufficiently represented by the selected empirical expressions. No claim of biophysical decoupling is made on the basis of selection frequency alone.

The MGGP algorithm generates a set of symbolic models with different predictive accuracy and complexity trade-offs, known as the Pareto set. The ten best candidate models were derived from one evolutionary run. To showcase the model selection procedure and candidates' diversity, we present the five symbolic equations highest in rank through Table 4. In terms of prediction accuracy, structural complexity and engineering interpretability, those five symbolic equations were screened for final structural model selection incorporated in the proposed GADT. At first, all the candidate variables given in Table 2 were considered as inputs. Following 100 independent MGGP runs, the variables having the highest frequencies and greatest consistency in occurrence across the winning symbolic equations are retained as final structural input set as detailed in Table 3.

Table 4 Top five MGGP candidate equations generated from a representative evolutionary run

Rank	Symbolic Equation *	Complexity	R^2
1	$y = 4.3(x_6) - 0.651(x_3) - 1.59e + 5\,(x_9)^3(x_{10})^6 + 4270.0(x_9)(x_{10})^2 0.326(x_6)^2 - 216.0(x_9)^{\frac{1}{2}} + 0.237(x_2)(x_6)(x_9)(x_{10})^2 + 64.3$	54	0.858
2	$y = 4.3(x_6) - 0.651(x_3) - 0.326(x_7) - 0.327(x_9) - 1.58e + 5\,(x_9)^3(x_{10})^6 + 4270.0(x_9)(x_{10})^2 - 0.326(x_6)^2 - 215.0(x_9)^{\frac{1}{2}} + 0.237(x_2)(x_6)(x_9)(x_{10})^2 + 66.3$	66	0.854
3	$y = 4.24(x_6) - 0.6(x_3) - 1.28e + 5\,(x_9)^3(x_{10})^6 + 4250.0(x_9)(x_{10})^2 - 0.32(x_6)^2 - 213.0(x_9)^{\frac{1}{2}} + 0.046(x_2)(x_6)(x_9)(x_{10}) + 57.3$	45	0.852
4	$y = 4010.0(x_9)(x_{10})^2 - 1.46e + 5\,(x_9)^3(x_{10})^6 - 3.62(x_6) - 204.0(x_9)^{\frac{1}{2}} + 0.172(x_2)(x_6)(x_9)(x_{10})^2 + 108$	40	0.831
5	$y = 2.94(x_6) - 0.459(x_3) - 1.16e + 5\,(x_9)^3(x_{10})^6 + 3910.0(x_9)(x_{10})^2 - 0.229(x_6)^2 - 188.0(x_9)^{\frac{1}{2}} + 55.8$	39	0.829

- Where $x_1, x_2, x_3, \ldots, x_{12}$ are input variables and output variable y as per table 2

The Complexity column reports the symbolic-complexity measure used during model evaluation. In this study, complexity is calculated as the sum of the node counts of all subtrees within a symbolic expression, following the expressional-complexity definition implemented in GPTIPS 2. For multigene individuals, the complexity values of the constituent genes are summed. Lower values indicate a more compact symbolic expression under this definition. This measure is used to compare expression size and should not be interpreted as a direct measure of predictive accuracy or physical validity.

Although all five models achieved comparable predictive performance, the final structural model was selected based on the combined criteria of prediction accuracy, symbolic compactness, robustness, and engineering interpretability. Among these candidate equations, Rank 1 was selected as the final structural model based on the combined criteria of predictive accuracy, symbolic compactness, robustness, and engineering interpretability

The final symbolic expression used to generate the structural prediction is

$$y = 4.3(x_6) - 0.651(x_3) - 1.59e + 5\,(x_9)^3(x_{10})^6 + 4270.0(x_9)(x_{10})^2 0.326(x_6)^2 - 216.0(x_9)^{\frac{1}{2}} + 0.237(x_2)(x_6)(x_9)(x_{10})^2 + 64.3$$

This equation remained invariant throughout the online synchronization stage. No coefficients, genes, or structural terms of the MGGP model were modified during GADT operation; only the low-dimensional gain and bias parameters were updated recursively.

3.5 Final Structural MGGP Model

The final structural MGGP model was developed using the robustly selected process variables identified from the repeated MGGP analysis: CO_2 off-gas, dissolved oxygen, O_2 off-gas, PAA feed rate, sugar feed rate, substrate concentration, pH, and aeration rate. The resulting MGGP model is an explicit empirical symbolic expression that captures the main nonlinear penicillin-concentration trends in the development data and is amenable to an efficient evaluation during online operation. The algebraic explicitness of the expression does not mean that it reflects the underlying biochemical mechanisms, nor will it extrapolate reliably beyond the operating domain represented in the dataset. As an empirical symbolic-regression model, operational conditions different from the development data may deteriorate its predictive performance. A difference between measured concentration and predicted concentration can arise due to batch-to-batch variability, measurement bias, unmodelled disturbance or a limitation of this fitted expression. The current model is proposed for forecasting and synchronisation in the

considered domain; the symbolic form of the model should notbe viewed as a confirmed mechanistic kinetic model, as prior validation will be required for any ex-trapolation beyond that domain [15].

3.6 Role of the Structural Model within the Adaptive Digital Twin

The MGGP model is a structure layer that does not change in the governed adaptive digital twin suggested. During the online deployment process, the symbolic expressions and their coefficients remain unchanged. This maintains the interpretability, model traceability and process knowledge derived from historical industrial data. The governed synchronization layer outside the MGGP model provides a remedy for the misalignment between the structural prediction and the evolving physical process, discussed in Section 4. This makes the digital twin accommodating for operating variabilities and, therefore, does not cause permanent structural retraining. As a result, the structural model provides the knowledge base that guarantees stability, while the adaptive layer is only responsible for synchronizing the models during batch operation[17].

4. Adaptive Digital Twin Architecture

4.1 Architectural Design Principle

The Governed Adaptive Digital Twin (GADT) aims to maintain the alignment between the digital twin and the physical fermentation process as it changes, without frequently altering the previously validated structural model. In industrial biomanufacturing, long batch durations, biological variability, sparse quality measurements, and process drift can progressively degrade a fixed, non-adaptive model and make the distinction important. Nonetheless, online retraining or structural model updates quite often can amplify the computational burden and reduce interpretability, complications in validation, and weak lifecycle traceability. The suggested architecture thus decouples structural process modelling from operational synchronisation. The structural model is constant process knowledge extracted from historical industrial data, while the synchronization layer addresses deviations occurring in individual batches during operation. Adaptation is not posited to be a repeated redevelopment of models in this framework but to be a governed synchronization problem. The symbolic model's interpretation and validation status remain safe while enabling the digital twin to respond to changes in the process[18].

For clarity, the terms used in this study are operational definitions of the evaluated configurations. The fixed, non-adaptive digital twin uses the offline MGGP expression without online parameter updates. The bias-adaptive digital twin updates an output-level bias correction while retaining the MGGP structure. The governed adaptive digital twin retains the same MGGP structure but updates both gain and bias through the specified recursive-estimation and governance procedure. Thus, “fixed” refers to the absence of online updates to the baseline model, not to an inability to predict time-varying fermentation trajectories[19].

A governed adaptive digital twin refers to a digital twin architecture, which means that a fixed structural process model is synchronised with the evolving physical process by means of constrained, interpretable, and explicitly governed online adaptation. Unlike conventional adaptive digital twins that may update model structures or high-dimensional parameter sets, only synchronization parameters required for agreement of measured and predicted process behaviour will be permitted to evolve during operation in the proposed GADT. This design approach is structured around five practical conditions for an industrial implementation: structural invariance; low adaptive complexity; clear operational governance; engineering intelligibility; and computational efficiency. Structural invariance guarantees that the symbolic MGGP model remains unchanged after offline validation. Low-dimensional parameter space limits online adaptation due to minimal adaptive complexity. Operational governance influences the timing and manner of adaptation. Engineering interpretability ensures adaptive corrections can be understood by process engineers. Efficiency avoids applying model-influenced control actions by relying on process data observed globally[19].

4.2 Overall Architecture of the Governed Adaptive Digital Twin

As depicted in (Fig. 2) the proposed GADT integrates five functional layers namely the physical fermentation process, data acquisition layer, structural digital twin, governed synchronization layer, and governance and decision layer. The physical process layer represents the fermentation that makes penicillin in fed-batch mode. The process trajectory for batches may differ, even if the operating recipes are similar. This is due to variances in the quality of inoculum, raw materials, oxygen transfer, and microbial activity. Biological factors can also result in a difference in trajectory. The evolving process forms the reference system that the digital twin must represent during the full batch. The routine measured process variables are a fermentation time, aeration rate, broth volume, base addition, sugar feed rate and phenyl acetic acid feed rate. This is what the data acquisition layer collects. The inputs of a structural digital twin are those that are constantly or regularly obtainable. On the contrary, online measurement of penicillin concentration does not take place continuously and, it typically becomes available through laboratory or pseudo-online quality testing conducted at regular intervals. As a result, the architecture assumes quality information in sequence and not perfect continuous measurements on products.

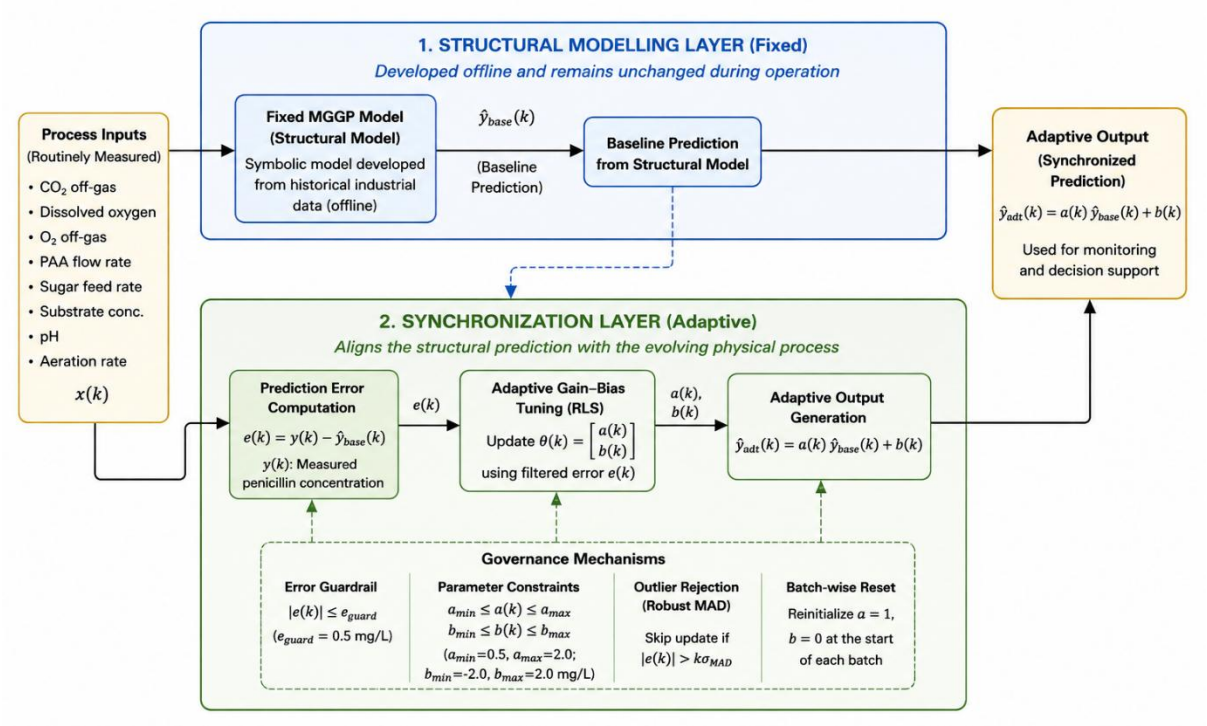

Fig. 2 Separation of the structural modelling layer and synchronization layer

The separation of the structural modelling layer and synchronization layer in the proposed Governed Adaptive Digital Twin (GADT) is illustrated in (Fig. 2). The structural layer keeps invariant process knowledge and the synchronization layer mitigates operational variability through governed online adaptation. The structural digital twin layer comprises an offline developed symbolic MGGP model presented in Section 3 that supports the baseline prediction of the penicillin concentration from available process variables. While operations are online, neither the symbolic equation nor its coefficients are changed. Accordingly, the MGGP model serves as the permanent knowledge layer of the digital twin that holds the nonlinear process relationships from historical industrial data. The governed synchronization layer operates on the output of the structural model rather than modifying the internal symbolic model. It serves to rectify the disparity between the structural forecast and the observed process state through a low-dimensional adaptive transformation. Synchronization is attained by updating two parameters in the proposed implementation. One is an adaptive gain and the other is an adaptive bias. The gain compensates for multiplicative variations in process performance or profile size, and the bias compensates for additive offset caused by sensor bias, unmodelled disturbance and residual error accumulation. Because only two parameters are updated, the online synchronization remains cheap in terms of computation and physically interpretable. The adaptive behaviour of synchronization layer is governed by governance layer and decision layer. The objective is not simply to improve prediction accuracy, but to preserve the physical meaningfulness, numerical stability, and industrial applicability of the adaptation. The regulation of adaptation through scheduling, innovation filtering, parameter bounds, and batch-wise reinitialization is performed by this layer.. The authorized prediction generated by the GADT can then be applied to the monitoring of a process, detecting deviations, operator support, or optimization of the future[20].

4.3 Governance Strategy and Architectural Implications

The suggested architecture is different from the conventional approaches of adaptive modelling as the synchronization and structural learning are treated as different functions. During a batch, not all prediction errors require a change in the symbolic model. These errors in many industrial fermentations are originate from a temporary or batch-related effect like physiological drift, raw material change, measurement uncertainty, or any operating disturbance. Thus, the synchronization layer merely corrects the operational mismatch while leaving the structural model intact in this framework[21].

This separation provides many benefits. First, model interpretability is preserved as the MGGP equation does not change at deployment. Secondly, the computational complexity is less since the online adaptation involves only two synchronization parameters and not the continuous optimization of a nonlinear symbolic model. In addition, the structural model remains unchanged across batches, simplifying validation and lifecycle management. The adaptive behaviour remains transparent, as every correction can be interpreted either as a scaling adjustment or an offset correction. In the end, separation defines a natural site for operational governance, allowing engineering constraints to be applied directly on the synchronization but without altering the structural model[21].

Table 5 outlines how traditional static digital twins, usual adaptive digital twins, and the GADT being proposed differ. Static digital twins maintain interpretability but may drift from the process. The conventional approach using adaptive digital twins achieves better accuracy in the short term. However, they either require the online retraining or high-dimensional model adaptation. This complicated furthering the validation as well as deployment. The proposed GADT consists of a fixed symbolic structural model combined with controlled low-dimensional synchronization. It achieves adaptability while preserving interpretability, computational performance and life cycle traceability.

This section gives a mathematical formulation of this architecture including synchronization model, recursive parameter estimator and governance rules employed for implementation on-line.

Table 5 Architectural comparison of conventional static digital twins, typical adaptive digital twins, and the proposed Governed Adaptive Digital Twin (GADT) for industrial bioprocess applications

Architectural Attribute	Conventional Static DT	Typical Adaptive DT	Proposed GADT
Structural model evolution	Fixed after offline development	Updated through online adaptation or periodic retraining	Fixed; structural model remains invariant
Online model retraining	Not required	Frequently required	Not required
Adaptation objective	Not applicable	Improve prediction accuracy through model evolution	Maintain long-term synchronization while preserving structural model integrity
Synchronization strategy	No explicit synchronization mechanism	High-dimensional parameter or model adaptation	Governed low-dimensional gain–bias adaptation
Model interpretability	High	Often reduced due to evolving model structure	High; symbolic structural model preserved
Adaptation governance	Not applicable	Generally limited or absent	Explicit scheduling, innovation filtering, parameter guardrails, and batch lifecycle management
Computational complexity	Low	Moderate to high	Low
Model lifecycle traceability	High	Reduced by continuous structural evolution	High; structural model remains unchanged and adaptation is fully traceable
Industrial deploy ability	Moderate; performance deteriorates under process drift	Challenging due to validation, maintenance, and computational overhead	High; stable, interpretable, computationally efficient, and suitable for long-term industrial deployment

5. Mathematical Formulation of the Governed Adaptive Digital Twin

5.1 Synchronization Objective

The objective of the Governed Adaptive Digital Twin (GADT) is to maintain synchronization between the physical fermentation process and its digital representation without modifying the validated structural MGGP model. At each discrete sampling instant, the routinely available process measurements are represented as

$$\mathbf{x}(k) \in \mathbb{R}^m$$

$$\mathbf{x}(k) = [x_{10}(k), x_6(k), x_{11}(k) x_3(k) x_2(k) x_{12}(k) x_8(k) x_4(k)]^T.$$

Here, (k) is the sampling index, and the process input vector contains the eight variables retained for the final structural digital twin: CO_2 off-gas, dissolved oxygen, O_2 off-gas, PAA floe rate, sugar feed rate, substrate concentration, pH, and aeration rate The fixed structural digital twin provides the symbolic prediction

$$\hat{y}_s(k) = f_{\mathrm{MGGP}}(\mathbf{x}(k)),$$

where $\hat{y}_s(k)$ is the structural MGGP prediction and $f_{\mathrm{MGGP}}(\cdot)$ is the invariant symbolic model obtained during offline training. The measured penicillin concentration is denoted by $y(k)$ and becomes available sequentially through pseudo-online quality measurements. The synchronization error is therefore defined as

$$e(k) = y(k) - \hat{y}_{\mathrm{ADT}}(k),$$

where $\hat{y}_{\mathrm{ADT}}(k)$ is the governed adaptive digital twin prediction. Thus, the online problem is not model redevelopment but the estimation of a low-dimensional correction that minimizes $e(k)$ while preserving the structural model.

5.2 Low-Dimensional Adaptive Synchronization Model

The synchronization layer acts only on the output of the structural digital twin. The adaptive prediction is expressed as

$$\hat{y}_{\mathrm{ADT}}(k) = \max[0,\ a(k)\hat{y}_s(k) + b(k)].$$

In this expression, $a(k)$ is the adaptive gain, $b(k)$ is the adaptive bias, and the maximum operator enforces non-negative penicillin predictions. The gain compensates for multiplicative changes in apparent productivity or trajectory magnitude, whereas the bias compensates for additive offsets caused by sensor bias, unmodelled disturbances, or accumulated residual error. Restricting adaptation to $a(k)$ and $b(k)$ preserves interpretability and prevents uncontrolled modification of $f_{\mathrm{MGGP}}(\cdot)$.

For recursive estimation, the adaptive parameter vector is written as

$$\boldsymbol{\theta}(k) = \begin{bmatrix} a(k) \\ b(k) \end{bmatrix},$$

with the regression vector

$$\boldsymbol{\phi}(k) = \begin{bmatrix} \hat{y}_s(k) \\ 1 \end{bmatrix}.$$

Before enforcing the non-negativity constraint, the affine synchronization model is linear in the adaptive parameters:

$$\hat{y}(k) = \boldsymbol{\phi}^T(k)\boldsymbol{\theta}(k).$$

Here, $\hat{y}(k)$ is the unconstrained synchronized prediction, $\boldsymbol{\phi}^T(k)$ is the transpose of the regression vector, and $\boldsymbol{\theta}(k)$ contains the two online synchronization parameters. This linear-in-parameters form enables efficient recursive estimation even though the underlying MGGP model is nonlinear.

The gain and bias are not biochemical kinetic parameters, they are empirical output-correction parameters. The gain alters the scale of the base prediction while the bias alters its ad∗dition. The objective of this low-dimensional adjustment is to account for systematic variations between the constant model output and the measured concentration while preserving the original MGGP expression. The observed improvement under the evaluation protocol motivated its use rather than an assumption of the linear nature of the underlying nonlinear fermentation dynamics. The affine approximation is not capable of representing all possible nonlinear model mismatches, nor do we know how it performs outside the validation domain.

5.3 Innovation Filtering and Recursive Parameter Estimation

Synchronization is driven by the innovation between the measured penicillin concentration and the unconstrained adaptive prediction:

$$e(k) = y(k) - [a(k)\hat{y}_s(k) + b(k)].$$

The innovation is computed before clipping the output to zero so that parameter estimation is not distorted when the displayed prediction reaches its physical lower bound. Because industrial quality measurements may contain laboratory variability, sampling uncertainty, or isolated abnormal observations, the raw innovation is filtered using an exponentially weighted moving average:

$$\overline{e}(k) = \rho e(k) + (1-\rho)\overline{e}(k-1), \qquad 0 < \rho < 1.$$

Here, $\overline{e}(k)$ is the filtered innovation, $\overline{e}(k-1)$ is its previous value, and ρ is the smoothing coefficient. Smaller ρ values provide stronger noise suppression, whereas larger values increase responsiveness to process changes.

The adaptive parameters are estimated with forgetting-factor recursive least squares, which updates the gain and bias from the prediction innovation with each new penicillin observation. The forgetting factor manages the influence of earlier observations and recent ones. The objective of adaptation uses the recursive estimation of the affine correction parameters as opposed to performing a separate batch-wide MAE-minimization routine at every sampling instant. The MAE and RMSE are used as evaluation measures unless the implementation decides otherwise.

The adaptive parameters are updated using forgetting-factor Recursive Least Squares (RLS). The estimator gain is calculated as

$$\mathbf{K}(k) = \frac{\mathbf{P}(k-1)\boldsymbol{\phi}(k)}{\lambda + \boldsymbol{\phi}^T(k)\mathbf{P}(k-1)\boldsymbol{\phi}(k)}, \qquad 0 < \lambda \leq 1,$$

where $\mathbf{K}(k)$ is the estimator gain, $\mathbf{P}(k-1)$ is the previous covariance matrix, and λ is the forgetting factor. The parameter vector is then updated using the filtered innovation:

$$\boldsymbol{\theta}(k) = \boldsymbol{\theta}(k-1) + \mathbf{K}(k)\overline{e}(k).$$

Finally, the covariance matrix evolves according to

$$\mathbf{P}(k) = \frac{1}{\lambda}[\mathbf{P}(k-1) - \mathbf{K}(k)\boldsymbol{\phi}^T(k)\mathbf{P}(k-1)].$$

When $\lambda = 1$, all previous observations contribute equally, producing conservative adaptation. When $\lambda < 1$, recent measurements receive greater weight, allowing the synchronization layer to follow slow physiological drift and batch-specific variability. Because only two parameters are estimated, the RLS update remains computationally inexpensive throughout long fermentation batches.

5.4 Governed Synchronization Mechanism

Recursive estimation is embedded within a governance framework to prevent unstable or physically unrealistic adaptation. First, scheduled adaptation prevents parameter updates during the initial growth phase, when penicillin concentration is close to zero and the innovation signal is weakly informative. Let t_{adapt} denote the adaptation activation time. Recursive synchronization is performed only when

$$t(k) \geq t_{\mathrm{adapt}}.$$

Before this condition is satisfied, the adaptive layer is initialized as

$$a(k) = 1, \qquad b(k) = 0,$$

so that the GADT output is identical to the structural MGGP prediction.

Second, an error-triggered responsiveness mechanism is used when the filtered innovation exceeds a predefined guard threshold:

$$|\bar{e}(k)| > e_{\mathrm{guard}}.$$

Here, e_{guard} is the threshold for activating enhanced synchronization. In this mode, the estimator temporarily becomes more responsive by reducing the effective forgetting factor and innovation smoothing. Once the filtered error returns below the threshold, normal conservative adaptation is restored.

Third, parameter projection constrains the adaptive gain and bias to admissible engineering ranges:

$$a_{\min} \leq a(k) \leq a_{\max},$$

$$b_{\min} \leq b(k) \leq b_{\max}.$$

The bounds $a_{\min}$, $a_{\max}$, $b_{\min}$, and $b_{\max}$ prevent numerical divergence and ensure that adaptive corrections remain physically meaningful.

Finally, lifecycle governance is imposed at the beginning of every batch:

$$a(0) = 1, \qquad b(0) = 0, \qquad \mathbf{P}(0) = \mathbf{P}_0.$$

So here, every batch starts from a validated structural model, whereas synchronization develops solely from the measurements taken during the batch. Abnormal adaptive states cannot be transferred between production batches. The full stepwise algorithmic representation of the governed synchronization procedure is presented in Supplementary (Fig. S1), while the closed-loop governed adaptation structure is displayed in Supplementary (Fig. S2).

5.5 Online Implementation and Computational Characteristics

At each sampling instant, the GADT evaluates the fixed MGGP model, computes the innovation, filters the innovation, updates the gain and bias by RLS when adaptation is permitted, applies the governance constraints, and stores the synchronized prediction. The structural model is evaluated directly from its symbolic equation, and no online retraining or nonlinear optimization is required. The synchronization layer estimates only the two components of $\boldsymbol{\theta}(k)$ using a 2×2 covariance matrix, while scheduling, filtering, parameter projection, guardrail evaluation, and batch-wise reinitialization require only scalar arithmetic and logical checks. The gain and bias are adaptive states updated at eligible measurement-update events during a batch; they are not fixed constants and are not independently prescribed functions of elapsed time. In the present pseudo-online evaluation, the penicillin observations are supplied sequentially at the stated sampling interval, so the update timing follows that protocol. In a real deployment, updates would be triggered by the availability of valid penicillin measurements, which may be delayed or irregular. The behaviour of the estimator under such measurement delays and gaps has not been independently validated in the present study. Consequently, the computational and memory requirements remain constant per sampling instant and do not increase with batch duration. In the present application, updates are executed at the same sampling interval as the available process measurements.

The online implementation workflow retained in the main manuscript is summarized in (Fig. 3). This (Fig. 3) shows the sequential execution of structural prediction, innovation processing, governance checks, recursive parameter updating, adaptive prediction generation, and batch-wise reinitialization. The resulting implementation characteristics are summarized in Table 6.The proposed formulation is computationally suitable for real-time implementation in supervisory monitoring or digital twin platforms; however, actual industrial deployment requires validation under online operating conditions.

Table 6 Real-time deployment: Computationally suitable for real-time implementation; requires validation under actual online operating conditions

Implementation Attribute	Proposed GADT
Structural model	Fixed symbolic MGGP model
Adaptive parameters	Gain and bias (2 parameters)
Online estimation method	Recursive Least Squares (RLS)
Adaptive parameter vector dimension	2
Covariance matrix dimension	2×2
Computational complexity	O(1) with respect to batch duration because the adaptive state dimension remains fixed at two parameters
Memory complexity	O(1)
Structural model retraining	Not required
Online optimization	Not required
Batch lifecycle management	Automatic parameter reset at batch initialization
Adaptation governance	Scheduling, innovation filtering, error-triggered responsiveness, parameter guardrails, and batch lifecycle management
Real-time deployment	Suitable for industrial implementation with minimal computational overhead

Practical tuning guidelines for the main synchronization hyperparameters are given in Table 7.

Table 7 Recommended Hyperparameters for the Proposed Governed Adaptive Digital Twin

Parameter	**Description**	**Typical Selection Guideline**	**Actual values used**	**Practical Role**
λ	Forgetting factor	Close to unity (e.g., 0.98–0.999)	0.999 (normal mode)-0.995 (fast/guardrail mode)	Balances adaptation speed and long-term stability
ρ	EWMA smoothing coefficient	Selected according to measurement noise	0.85 (normal mode)-0.70 (fast/guardrail mode)	Suppresses noise while preserving responsiveness
t_{adapt}	Adaptation activation time	Beginning of the penicillin production phase	25 h	Prevents learning during non-informative operating conditions
e_{guard}	Innovation guard threshold	Based on historical innovation statistics	0.5	Activates enhanced synchronization only when required
a_{min}, a_{max}	Adaptive gain limits	Engineering-defined	0.5,2	Maintains physically meaningful scaling corrections
b_{min}, b_{max}	Adaptive bias limits	Engineering-defined	-2.0,+2	Prevents unrealistic offset corrections
Δ_{max}	Maximum prediction rate change	Based on expected process dynamics	Not implemented	Ensures smooth biological trajectory evolution
P_0	Initial covariance matrix	Large positive diagonal matrix	1000×I2	Represents initial uncertainty in adaptive parameters

The proposed formulation is therefore suitable for real-time industrial implementation in supervisory monitoring or digital twin platforms, while preserving the validated symbolic process model and ensuring that online adaptation remains governed, interpretable, and traceable.

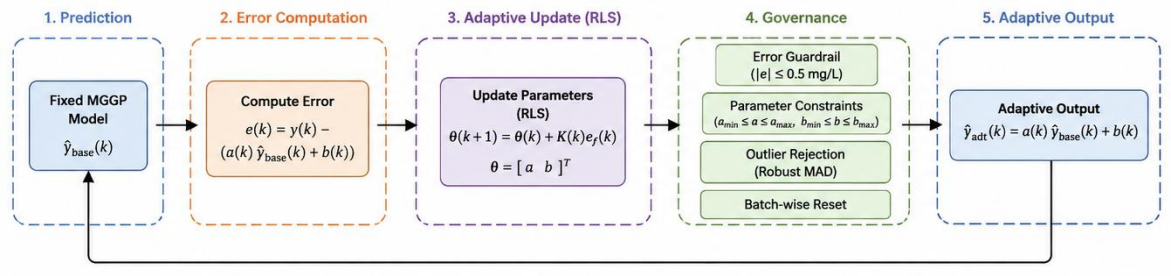

Fig. 3 Online implementation workflow of the proposed Governed Adaptive Digital Twin, showing sequential measurement acquisition, structural prediction, innovation processing, governed recursive synchronization, adaptive prediction generation, and batch-wise reinitialization

6. Experimental Design and Evaluation Protocol

6.1 Evaluation Framework

The Governed Adaptive Digital Twin (GADT) was tested using a pseudo-online streaming framework, which mimicked the working conditions of a penicillin fermentation industry. Validation that is direct under the live production conditions is rarely possible because the online measurement of the product is limited, the campaigns of production are costly, and are generally impractical for controlled experimentation on commercial bioprocesses. The historical campaign data were replayed sequentially according to the order used in the evaluation protocol. This procedure provides a stream of observations to the digital twin and enables pseudo-online assessment of the synchronization algorithm. The replay order is an evaluation procedure and should not, by itself, be interpreted as evidence that the campaigns represent a verified chronological production history or a measured process-drift trajectory. The evaluation dataset comprises of 90 simulated fed-batch fermentation campaigns from the industrial-scale IndPenSim benchmark which were collected over three years' time. With routinely measured process variables of sampling interval 0.2 h. Simulating the operation, the different batches were executed in order, and the measurements were presented to the digital twin one sample at a time, thereby replicating the information available during real operation. At each sampling instant, the symbolic MGGP model produced the structural prediction. The governed synchronization algorithm updated the adaptive gain and bias parameters whenever the adaptation conditions of Section 5 were satisfied.

Each fermentation batch was treated as an independent evaluation episode. At the start of every batch, the adaptive parameters and covariance matrix were reset to their nominal values, ensuring that the sync could only evolve based on information from the production run. The established protocols exhibit correspondence with the lifecycle governance strategy elaborated in Section 4 and also reflects industrial validation practices where adaptive states are not automatically passed on between independent production batches.

6.2 Experimental Protocol

To quantify the contribution of the proposed synchronization strategy, three digital twin configurations were evaluated under identical streaming conditions:

1. Static Digital Twin (SDT): the fixed symbolic MGGP model without online adaptation.

2. Bias-Adaptive Digital Twin (BADT): synchronization using adaptive bias correction only.

3. Governed Adaptive Digital Twin (GADT): the complete proposed framework incorporating adaptive gain–bias synchronization together with scheduling, recursive estimation, and governance mechanisms.

All models received identical process measurements and pseudo-online penicillin observations. Consequently, any performance differences can be attributed exclusively to the proposed synchronization strategy rather than to differences in process data or model structure.

To further assess the contribution of the governance framework, ablation studies were performed by selectively disabling individual synchronization components, including innovation filtering, adaptation scheduling, and parameter constraints. These experiments quantify the importance of each design element and demonstrate its contribution to stable long-term synchronization.

6.3 Performance Evaluation

Performance was evaluated using batch-wise prediction and synchronization metrics. Root Mean Square Error (RMSE) summarizes the typical magnitude of prediction error while giving greater weight to larger errors. Mean Absolute Error (MAE) summarizes the average absolute difference between measured and predicted penicillin concentrations. Batch-end absolute error measures the absolute prediction error at the final evaluated observation of each batch. Maximum trajectory absolute error denotes the largest absolute prediction error observed over the evaluated observations within a batch. The reported means and standard deviations summarize the distribution of each metric across the evaluated batches. These definitions are used consistently in the tables, results text, and figure captions.

Along with numerical performance assessment measures, synchronization behaviour was assessed with trajectory comparisons of measured penicillin concentration, static symbolic model and proffered adaptive digital twin. Great care was taken to see whether it was possible for GADT to compensate for gradual drift of the process and recover from accumulated prediction errorswhile showing stable adaptive behaviour in the presence of realistic trends whatever their origin is.

For each of the 90 fermentation batches, each performance metric was computed independently and then aggregated afterwards to understand robustness/generalisability across multiple production campaigns. A common evaluation protocol along with common process data will ensure a fair and reproducible comparison of synchronization performance of all competing models.

Results that are presented in the next section prove the efficiency of the proposed governed synchronization strategy that maintains long-term synchronization between the physical fermentation process and the digital twin under industry-representative operating conditions.

7. Results, Comparative Evaluation, and Discussion

7.1 Baseline Performance of the Static Digital Twin

The static symbolic digital twin was first evaluated to establish the baseline prediction performance before introducing online synchronization. (Fig. 4) compares representative measured penicillin concentration trajectories with the corresponding predictions generated by the fixed MGGP model. The fixed, non-adaptive model captures the main nonlinear trend of penicillin accumulation, including the onset of product formation and the subsequent production phase, confirming that the symbolic model provides a suitable structural representation of the fermentation process. However, the trajectories also show that prediction errors increase gradually as the batch progresses, especially during the later production stage. The batches shown in Fig. 4 were selected using a predefined, position-based illustrative selection criterion. Their ordinal IDs are Batch 3 and Batch 5, corresponding to the third and fifth consecutive fermentation trajectories, respectively, in the original dataset. These batches were selected to illustrate the digital twin's trajectory-prediction behaviour across two individual fermentation runs. They were not selected on the basis of median performance or statistical representativeness and should therefore be interpreted as illustrative examples rather than representative samples of the full 90-batch evaluation.

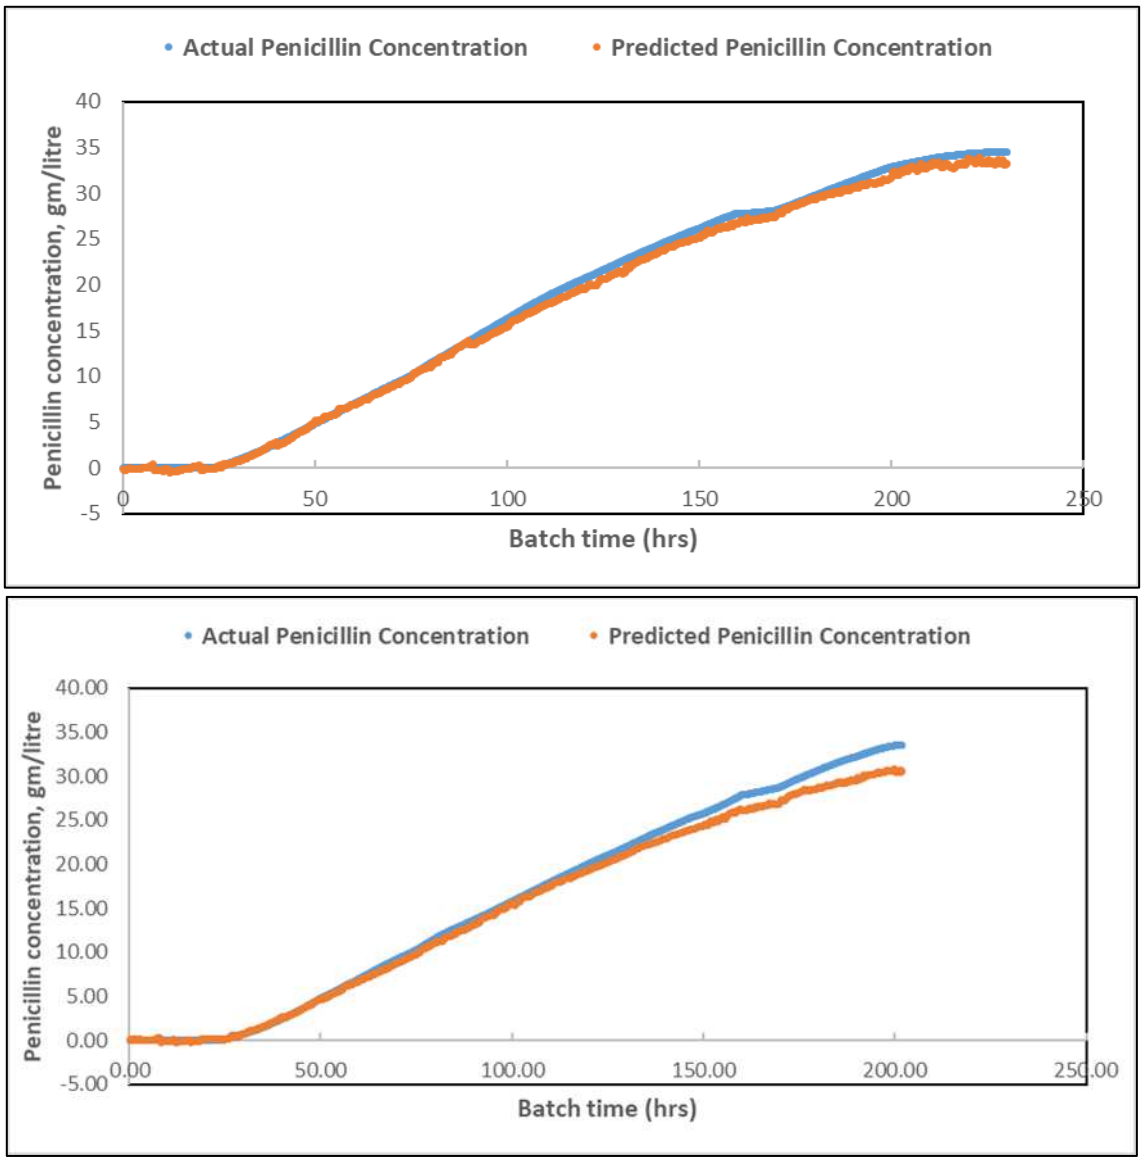

Fig. 4 Measured penicillin concentration and predictions from the fixed, non-adaptive MGGP model for Batch 3 and Batch 5. The displayed batches were selected using a predefined, position-based illustrative selection criterion, corresponding to the third and fifth consecutive fermentation trajectories in the original dataset. These trajectories illustrate the behaviour of the baseline model

in two individual fermentation runs and should be interpreted alongside the aggregate performance results across all 90 evaluated batches. This behaviour indicates that the main limitation of the fixed, non-adaptive model is not its inability to represent the overall fermentation trend, but its lack of an online mechanism to compensate for batch-specific drift. In long-duration fed-batch fermentation, small early deviations can accumulate into noticeable trajectory-level mismatch because microbial physiology, raw material effects, sensor bias, and operating disturbances may differ from batch to batch. Since the MGGP model remains fixed after offline development, it cannot adjust to these evolving conditions during deployment.

The quantitative baseline results over all 90 batches of IndPenSim industrial-scale penicillin fermentation benchmark are summarized in Table 8. The static digital twin achieved a mean RMSE of **0.67 ± 0.26 g/L**, indicating acceptable average predictive capability. However, the average batch-end absolute error was **0.88 ± 0.62 g/L**, and the maximum trajectory deviation reached **1.65 ± 0.51 g/L**. These larger trajectory-level errors confirm that synchronization loss becomes more pronounced over extended batch operation, even when the overall model structure remains valid.

Table 8 Baseline predictive performance of the static symbolic digital twin evaluated over 90 simulated fed-batch fermentation campaigns from the industrial-scale IndPenSim benchmark

Metric	Mean	Standard Deviation
RMSE (g/L)	**0.67**	**0.26**
Batch-end absolute error (g/L)	**0.88**	**0.62**
Maximum trajectory deviation (g/L)	**1.65**	**0.51**

Together, (Fig. 4) and Table 8 show that the fixed MGGP model is an appropriate structural foundation but is insufficient as a standalone digital twin for long-horizon industrial deployment. The results therefore justify the proposed GADT strategy: instead of retraining or modifying the symbolic model, a governed synchronization layer is required to correct gradual batch-specific deviations while preserving the validated structural model.

7.2 Single-Batch Demonstration of Governed Synchronization

To illustrate the operating behaviour of the proposed Governed Adaptive Digital Twin (GADT), a representative industrial fermentation batch was examined in detail. This analysis focuses on trajectory-level synchronization rather than only pointwise prediction accuracy. (Fig. 5a) compares the measured penicillin concentration with predictions from the static MGGP digital twin and the proposed GADT. During the initial growth phase, penicillin concentration remains close to zero and both models generate similar predictions because adaptation is intentionally inactive before informative product measurements become available. This confirms that the scheduled adaptation rule prevents unnecessary parameter updates during the low-information stage of fermentation.

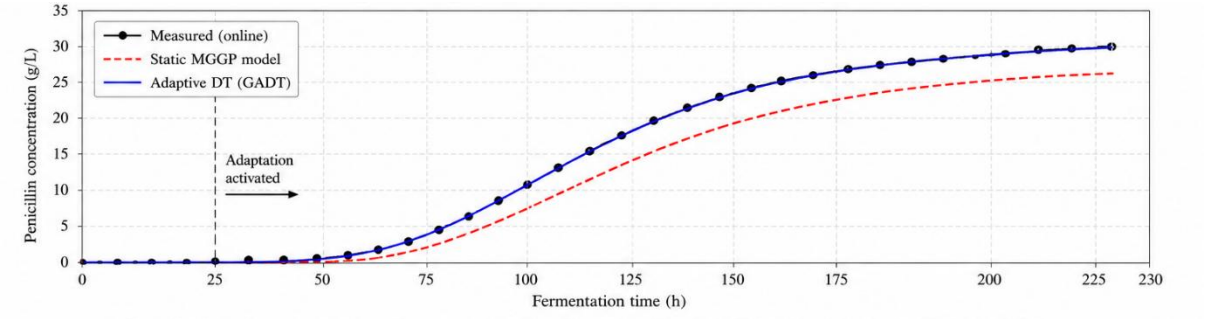

Fig. 5a Representative Batch: Measured and Predicted Penicillin Concentration Trajectories

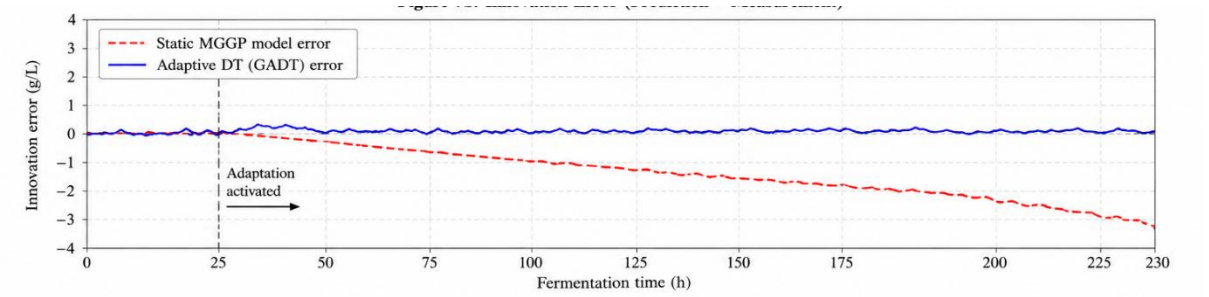

Fig. 5b Innovation Error Profiles of the Static MGGP Model and the Governed Adaptive Digital Twin (GADT)

7.3 Multi-Batch Evaluation of Adaptive Digital Twin Performance

The proposed GADT was next evaluated against the static digital twin and the bias-adaptive digital twin over the complete dataset of 90 simulated fed-batch fermentation campaigns generated using the IndPenSim industrial-scale penicillin fermentation simulator, labelled as representing approximately three years of simulated production campaigns in the IndPenSim benchmark. Table 9 summarizes the aggregate prediction and synchronization performance of the three configurations.

Table 9 Comparative prediction and synchronization performance of the static digital twin, bias-adaptive digital twin, and proposed GADT over 90 simulated fermentation batches derived from the IndPenSim industrial-scale penicillin fermentation benchmark.

Improvement over Static DT	BADT	GADT
RMSE reduction	**39%**	**46%**
MAE reduction	**40%**	**49%**
Batch-end error reduction	**38%**	**68%**
Maximum trajectory deviation reduction	**27%**	**70%**

Across all evaluation metrics, the proposed GADT provides the best performance. The mean RMSE is reduced from **0.67 ± 0.26 g/L** for the static digital twin to **0.36 ± 0.14 g/L** for the GADT, corresponding to an improvement of approximately **46%**. The mean absolute error also decreases from **0.53 ± 0.21 g/L** to **0.27 ± 0.11 g/L**, confirming that the proposed synchronization layer improves the overall trajectory agreement rather than only isolated prediction points. The improvement is more pronounced for long-horizon synchronization metrics. The average batch-end absolute error decreases from **0.88 ± 0.62 g/L** to **0.28 ± 0.18 g/L**, representing an improvement of nearly **68%**, while the maximum trajectory deviation decreases from **1.65 ± 0.51 g/L** to **0.49 ± 0.19 g/L**, corresponding to an improvement of approximately 70%. The observed improvement indicates that governed gain–bias adaptation reduces prediction mismatch under the pseudo-online evaluation protocol. The results demonstrate improved agreement between predicted and measured penicillin concentrations across the evaluated batches. However, the present evaluation does not independently identify the physical causes of the baseline prediction errors or establish that they arise specifically from chronological process drift. The gradual correction of prediction mismatch is consistent with the intended role of the synchronization layer, while the relative contribution of biological variability, measurement effects, and other unmodelled factors remains unresolved.

This creates a confirming evidence of the stabilising effect of governed synchronisation in Fig. 5(b). The static digital twin shows a persistently growing innovation pattern because of the accumulation of trajectory mismatch, while the GADT reduces the magnitude of the innovation and keeps the error more closely around zero once adapted. The behaviour shows that the adaptive layer does not replace the symbolic one, but instead compensates for the gradual mismatch between the fixed structural prediction and the physical process.

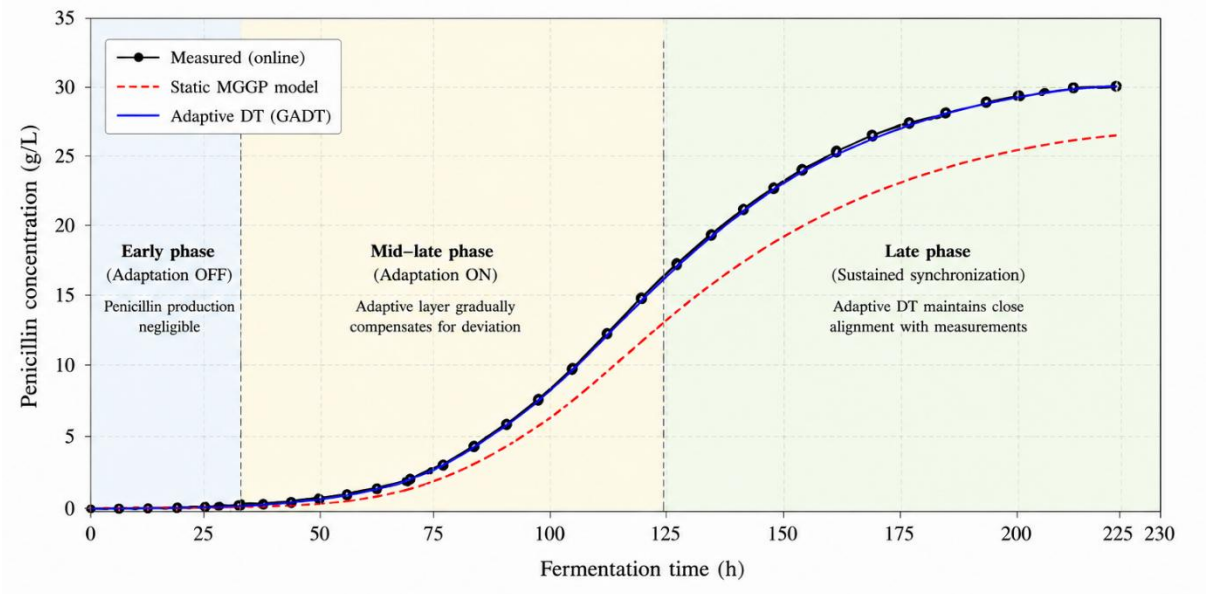

Fig. 6 Synchronization trajectories from the proposed Governed Adaptive Digital Twin (GADT)

Fig. 6 shows the penicillin measurements, associated static MGGP predictions, and governed adaptive predictions over a long-horizon fermentation. As shown in Table 9 and (Fig. 6) having realistic pseudo-online industrial conditions, the predictions of average accuracy, and trajectory level synchronization of GADT has improved. First, these improvements are achieved without changing or retraining the structural

MGGP model and thus support the main claim of this paper that we can improve long-term reliability of digital twins through governed synchronization rather than through continuous structural model evolution.

7.4 Generalization Across Campaigns and Long-Horizon Synchronization

To assess the consistency of the synchronization strategy across the 90 evaluated fermentation campaigns, batch-wise performance was examined alongside the aggregate results reported in Section 7.3. This analysis evaluates whether the performance advantage of the governed adaptive digital twin is maintained across the evaluation sequence. Because the campaign labels and replay order do not independently establish chronological process drift, the results are interpreted as evidence of performance across the evaluated campaigns rather than as proof of a specific temporal drift mechanism. The static MGGP model provides a stable structural representation of the fermentation process but exhibits progressively larger prediction errors as operating conditions deviate from those represented in the historical training data. Introducing adaptive bias correction improves prediction accuracy by reducing systematic offsets; however, the proposed GADT consistently achieves the best overall performance because simultaneous gain and bias adaptation compensates for both additive and multiplicative deviations while preserving the fixed symbolic model.

Long-term digital twin deployment requires stable performance not only within individual batches but also across extended production campaigns. To assess this capability, the 90 fermentation batches were arranged chronologically according to production sequence, and the batch-wise RMSE was calculated for both the static MGGP digital twin and the proposed GADT. The resulting long-horizon comparison is shown in (Fig. 7).

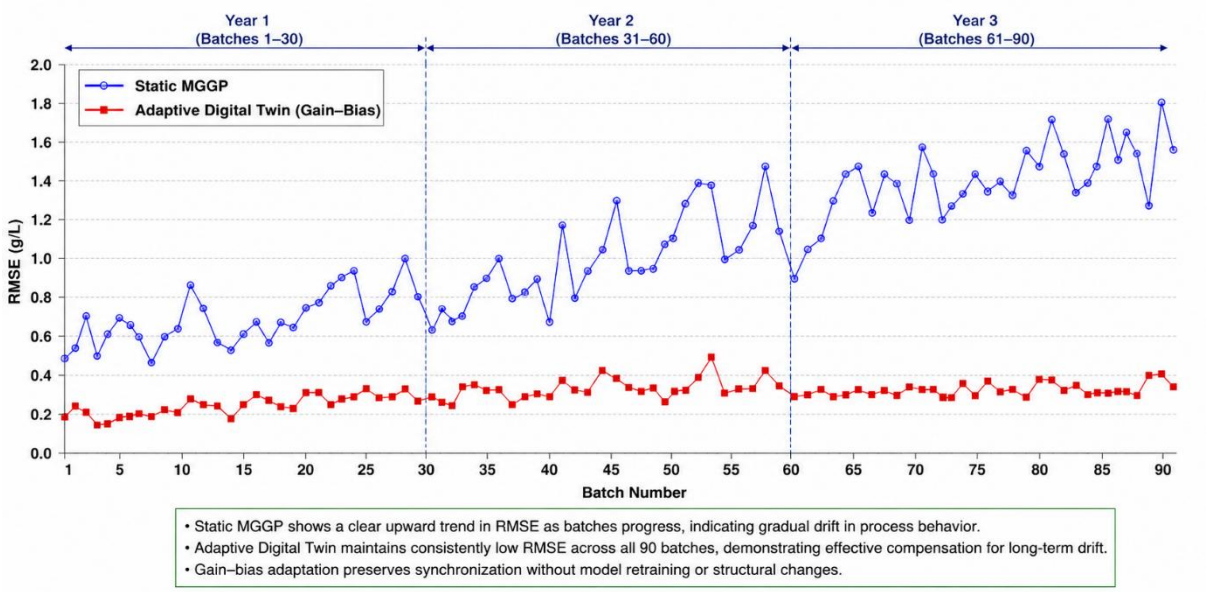

Fig. 7 Long-horizon batch-wise RMSE across 90 simulated fermentation batches derived from the IndPenSim industrial-scale penicillin fermentation benchmark, comparing the static MGGP digital twin and the proposed Governed Adaptive Digital Twin (GADT). The GADT maintains more consistent prediction accuracy across the production campaign by compensating for gradual process drift through governed synchronization

The static MGGP model shows increasing prediction error at intervals, indicating that with process operation this process has lost synchronization with the conditions when the model was developed offline. The behaviour is also expected during long-duration biomanufacturing campaigns, where gradual changes over time may occur in inoculum quality, raw material variability, microbial physiology, equipment condition and measurement characteristics. Because the static symbolic model remains unchanged after offline development, it cannot explicitly compensate for these gradual process changes during deployment. Conversely, the GADT proposal shows a less varied prediction accuracy throughout the whole 3 years' data. The governed synchronization layer keeps adjusting the gain and the bias as new quality measurements are available. This helps to counteract slow process drift before large trajectory-level errors build up. We can confirm that the improvement achieved in subsections 5.1 to 5.3 is not merely due to representative batches, but a full industrial production.

In summation, (Fig. 7) shows that the proposed framework can maintain synchronization despite realistic process drift, improving long-horizon robustness. As per these results, the main idea of this work is supported, which indicates that the reliability of a digital twin for the MGGP model is achievable by regulated low-dimensional synchronisation without the need of repeated structural retraining.

7.5 Engineering Role of the Governance Mechanisms

The distinguishing feature of the proposed framework is not adaptation alone but the explicit governance of adaptation. To examine the contribution of individual governance components, an ablation study was conducted on the representative batch used for the single-batch demonstration. Adaptation scheduling, parameter guardrails, and innovation filtering were selectively disabled while the remaining synchronization algorithm was kept unchanged. The corresponding results are summarized in Table 10. Because this ablation analysis uses the representative batch, its absolute error values are not directly comparable with the 90-batch aggregate results reported in Table 9.

Table 10 Ablation analysis of the governance mechanisms in the proposed GADT

Configuration	RMSE	MaxError	ObservedBehavior
Full GADT	1.10	5.52	Stable, bounded
No scheduling	0.89	4.15	Early-batch instability
No guardrails	1.14	6.52	Occasional overshoot
No error filtering	0.89	4.15	Increased noise sensitivity

According to the ablation study, using prediction accuracy as an evaluation of adaptive digital twins is not enough. A number of reduced configurations yield average values of RMSE that are similar to those of the complete framework; but this misleading result is due to unstable adaptive behaviour, rather than improved synchronisation. By taking away the scheduling of adaptation we allow the parameters to be updated at the early times of fermentation when the penicillin concentration is practically zero and the synchronization errors are weakly informative. In such a scenario, an adaptive estimator reacts excessively to measurement uncertainty which leads to unstable evolution of parameters, onwards cascading throughout the rest of the batch.

Likewise, the local disturbances and/or noise are allowed to update the parameters freely without guardrails. While the average prediction error does not vary a lot, the maximum prediction error does vary quite a lot. This indicates that occasionally synchronization is lost. Such behaviour would not be acceptable in the industrial context, where a reliable operation is generally far more important than slight gains in average statistical performance. Turning off error filtering causes a similar degradation because transitory measurement fluctuations are directly impacting estimation of the parameters. The synchronization that results becomes more sensitive to measurement noise and numerically less robust with prolonged fermentation. As a case in point, governance mechanisms must not be regarded as implementation detail. Instead, they must be considered an architectural feature of adaptive digital twins. The suggested framework focuses on stability, interpretability, and operational reliability instead of maximizing adaptive flexibility due to the practical demands of industrial biomanufacturing.

7.6 Engineering Interpretation of Adaptive Synchronization Behaviour

The results indicate that gain–bias adaptation can reduce the prediction mismatch observed between the fixed MGGP model and measured penicillin concentrations under the evaluated conditions. The observed improvement is consistent with the use of a low-dimensional output correction to address scale and offset discrepancies. However, the current analysis does not establish that all prediction errors arise from scale and offset deviations, nor does it identify their underlying biochemical or operational causes. This provides an empirical rationale for using two output-correction parameters in the evaluated setting. The gain adjusts prediction scale and the bias adjusts prediction offset, but their effectiveness should not be interpreted as evidence that the underlying fermentation dynamics are governed by only two physical mechanisms.

Together with the ablation results summarized in Table 10, this behaviour confirms that the improved performance of the GADT is not the result of unconstrained online learning. Rather, synchronization is achieved through bounded, interpretable, and governed correction of the structural prediction. This finding is important for industrial use because it shows that long-term digital twin alignment can be maintained without continuously modifying the validated symbolic model.

7.7 Implications for Industrial Digital Twin Deployment

The proposed GADT addresses a key practical limitation of adaptive digital twins in biomanufacturing: the need to balance adaptability with interpretability, validation, and lifecycle traceability. Many adaptive modelling approaches rely on high-dimensional parameter updating, online retraining, or evolving machine-learning structures. Although such methods can improve short-term prediction accuracy, they may increase computational cost, reduce transparency, and complicate model verification during deployment.

In contrast, the proposed framework restricts online adaptation to a governed synchronization layer acting on only gain and bias parameters. The MGGP model remains fixed throughout operation, preserving the process knowledge extracted from historical industrial data. This architecture simplifies validation because the structural model does not change, while the adaptive behaviour remains easy to interpret as either a scaling correction or an offset correction.

From an implementation perspective, the low-dimensional recursive update is computationally light and compatible with existing supervisory monitoring systems. The governance mechanisms further improve industrial suitability by ensuring that adaptation is scheduled, bounded, filtered, and reset at the batch level. Therefore, the results suggest that adaptive digital twins for regulated biomanufacturing should be designed as governed synchronization systems rather than continuously retrained predictive models.

7.8 Limitations and Future Research

Several limitations should be acknowledged. First, although the study used 90 simulated fed-batch fermentation campaigns from the industrial-scale IndPenSim benchmark over approximately three years, synchronization was evaluated under pseudo-online conditions using historical quality measurements. Future work should evaluate the framework under actual online deployment, including realistic laboratory delays, irregular sampling, missing measurements, and measurement uncertainty.

The reported results summarize point-prediction performance and variability in the evaluation metrics across batches; they do not constitute formal prediction intervals or calibrated confidence bounds for individual trajectories. Consequently, the present results do not quantify the

uncertainty associated with each predicted penicillin concentration. Future work should evaluate residual-based uncertainty estimation or another appropriate interval-calibration method and assess its coverage under delayed, irregular, and missing quality measurements.

Second, the present synchronization layer is intentionally limited to gain and bias adaptation. This low-dimensional formulation was effective for the industrial dataset studied here, but more severe process changes, equipment faults, or major shifts in microbial physiology may require additional adaptive degrees of freedom or hybrid synchronization strategies. Future studies should examine when gain–bias adaptation is sufficient and when richer adaptive structures become necessary.

Third, the present work focuses on monitoring and soft-sensing. Extending governed synchronization to model predictive control, real-time optimization, fault diagnosis, and operator decision support represents an important direction for future research. Finally, the framework should be validated on other batch and continuous bioprocesses to confirm the broader generality of the governed synchronization concept beyond penicillin fermentation.

8. Conclusions

This study proposed a Governed Adaptive Digital Twin (GADT) framework for maintaining long-term synchronization in industrial penicillin fermentation. The framework separates structural process modelling from online synchronization: a symbolic MGGP model provides the fixed structural representation of the process, while a low-dimensional adaptive layer updates only gain and bias parameters to correct batch-specific deviations. This design preserves model interpretability and lifecycle traceability while avoiding repeated structural retraining. The proposed framework was evaluated using pseudo-online replay of 90 simulated fed-batch fermentation campaigns from the IndPenSim industrial-scale benchmark, labelled as representing approximately three years of simulated production campaigns in the IndPenSim benchmark. Compared with the static digital twin, the GADT reduced RMSE from **0.67 ± 0.26 g/L** to **0.36 ± 0.14 g/L**, batch-end error from **0.88 ± 0.62 g/L** to **0.28 ± 0.18 g/L**, and maximum trajectory deviation from **1.65 ± 0.51 g/L** to **0.49 ± 0.19 g/L**. These results demonstrate that governed gain–bias synchronization can substantially reduce prediction error across the evaluated campaigns under the specified pseudo-online protocol, while preserving the underlying MGGP model structure. The evaluation supports improved synchronization performance but does not independently establish chronological process drift as the cause of the observed prediction mismatch.

The implemented governance mechanisms, including scheduled adaptation, innovation filtering, parameter constraints, and batch-wise reinitialization, were essential for stable and physically meaningful adaptation. Rate limiting was considered as an additional safeguard for practical deployment but was not implemented in the present study.. Overall, the results demonstrate that adaptive digital twins for industrial biomanufacturing can be effectively implemented as governed synchronization systems rather than continuously retrained predictive models. Future work should evaluate the framework under real online deployment with measurement delays, missing samples, and extension to control, optimization, and other bioprocesses.

Conflict of interest

The authors declare that they have not known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors

Author contributions

- Suchandan Banerjee: Methodology, Software, Data curation, Formal analysis, Investigation, Visualization, Writing – original draft
- Shrestha Giri: Data curation, Investigation
- Mithu Paul: Data curation, Investigation,
- Sandip Kumar Lahiri: Conceptualization, Supervision, Validation, Writing – review & editing.

Data availability

The data underlying this study are publicly available from Mendeley Data: “Data for: Modern day monitoring and control challenges outlined on an industrial-scale benchmark fermentation process,” Version 2, DOI: 10.17632/pdnjz7zz5x.2. The present study uses the 90 normal fermentation batches from the publicly available IndPenSim benchmark dataset.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

The authors acknowledge that portions of the manuscript, particularly in improving the clarity of language and structure, were supported using AI-assisted tools (e.g., ChatGPT by OpenAI). However, all scientific content, analyses, and interpretations were developed, reviewed, and verified by the authors. The authors take full responsibility for the accuracy, originality, and integrity of the information presented in this paper.

References

1. Udugama IA, Lopez PC, Gargalo CL, et al (2021) Digital Twin in biomanufacturing: challenges and opportunities towards its implementation. Syst Microbiol and Biomanuf 1:257–274. https://doi.org/10.1007/s43393-021-00024-0

2. Liu C, Gong Z, Feng E, Yin H (2012) Nonlinear dynamical systems of fed-batch fermentation and their optimal control. International Journal of Systems Science 43:809–819. https://doi.org/10.1080/00207721.2010.523800

3. Wang J, Moreira J, Cao Y, Gopaluni RB (2023) Simultaneous digital twin identification and signal-noise decomposition through modified generalized sparse identification of nonlinear dynamics. Computers & Chemical Engineering 177:108294. https://doi.org/10.1016/j.compchemeng.2023.108294

4. Moser A, Appl C, Pörtner R, et al (2024) A New Concept for the Rapid Development of Digital Twin Core Models for Bioprocesses in Various Reactor Designs. Fermentation 10:463. https://doi.org/10.3390/fermentation10090463

5. Kadlec P (2008) Soft sensors in bioprocessing: A review. Analytica Chimica Acta 621:1–17

6. Helgers H, Hengelbrock A, Rosengarten JF, et al (2022) Towards Autonomous Process Control—Digital Twin for HIV-Gag VLP Production in HEK293 Cells Using a Dynamic Metabolic Model. Processes 10:2015. https://doi.org/10.3390/pr10102015

7. Noh H, Kwon S, Seo IW, et al (2020) Multi-Gene Genetic Programming Regression Model for Prediction of Transient Storage Model Parameters in Natural Rivers. Water 13:76. https://doi.org/10.3390/w13010076

8. Yan X, Mohammadian A (2019) Multigene Genetic-Programming-Based Models for Initial Dilution of Laterally Confined Vertical Buoyant Jets. JMSE 7:246. https://doi.org/10.3390/jmse7080246

9. Goldrick S, Duran-Villalobos CA, Jankauskas K, et al (2019) Modern day monitoring and control challenges outlined on an industrial-scale benchmark fermentation process. Computers & Chemical Engineering 130:106471. https://doi.org/10.1016/j.compchemeng.2019.05.037

10. Goldrick S, Ştefan A, Lovett D, et al (2015) The development of an industrial-scale fed-batch fermentation simulation. Journal of Biotechnology 193:70–82. https://doi.org/10.1016/j.jbiotec.2014.10.029

11. Vaitkus V, Brazauskas K, Repšytė J (2020) Soft-sensors based on Black-box Models for Bioreactors Monitoring and State Estimation. In: Proceedings of the 2020 12th International Conference on Bioinformatics and Biomedical Technology. ACM, Xi'an China, pp 43–47

12. Hassan M, Montague G, Iqbal MZ, Fahey J (2024) Virtual reality-based bioreactor digital twin for operator training. Digital Chemical Engineering 11:100147. https://doi.org/10.1016/j.dche.2024.100147

13. Walmsley TG, Patros P, Yu W, et al (2024) Adaptive digital twins for energy-intensive industries and their local communities. Digital Chemical Engineering 10:100139. https://doi.org/10.1016/j.dche.2024.100139

14. Niu C, Li B, Wang Z (2023) Using artificial intelligence-based algorithms to identify critical fouling factors and predict fouling behavior in anaerobic membrane bioreactors. Journal of Membrane Science 687:122076. https://doi.org/10.1016/j.memsci.2023.122076

15. Kusznir T, Smoczek J (2022) Multi-Gene Genetic Programming-Based Identification of a Dynamic Prediction Model of an Overhead Traveling Crane. Sensors (Basel) 22:339. https://doi.org/10.3390/s22010339

16. Chakraborty A, Sivaram A, Samavedham L, Venkatasubramanian V (2020) Mechanism discovery and model identification using genetic feature extraction and statistical testing. Computers & Chemical Engineering 140:106900. https://doi.org/10.1016/j.compchemeng.2020.106900

17. Alghamdi W, Albassam E (2024) Synchronization Patterns for Digital Twin Systems. Journal of Applied Data Sciences 5:1026–1037. https://doi.org/10.47738/jads.v5i3.267

18. Stegagno P, Rossi M, Biagi P (2020) Challenges in adaptive digital twin design. IEEE Access 8:117030–117043

19. Ogunsakin R, Mehandjiev N, Marin CA (2023) Towards adaptive digital twins architecture. Computers in Industry 149:103920. https://doi.org/10.1016/j.compind.2023.103920

20. Erkoyuncu JA, del Amo IF, Ariansyah D, et al (2020) A design framework for adaptive digital twins. CIRP Annals 69:145–148. https://doi.org/10.1016/j.cirp.2020.04.086

21. Wang G, Zhao L (2020) Synchronization challenges in real-time bioprocess models. Computers & Chemical Engineering 135:106742

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryFigures.docx