Practical Advances in Artificial Intelligence and Computing
Frontier Research Publishing • https://www.frpubs.org
Vol. 4, No. 4, 2025

# When Online Growth Adaptation Helps a Synthetic Continuous Fermentation

**Olivren Zephwick[1]*, Ceryndra Valequinn[1], Thamiel Rooksen[1]**
[1]University of Guelph, Guelph, Canada
*Corresponding author: olivrenzephwick185@outlook.com

**Abstract:** Online parameter adaptation and biomass protection can alter a continuous fermentation differently during steady operation, activity loss and startup. We execute a two-state Monod chemostat simulation with 24 independently fitted kinetic models, five conditions and four dilution policies, giving 480 eighty-hour episodes. The policies are a frozen steady optimum, an online adaptive steady optimum, a protected adaptive rule and a truth-informed protected reference. No predictive controller, physical fermentation or organism-specific dataset is claimed. After a maximum-growth drop, mean cumulative biomass harvest is 25.54 for the frozen rule, 34.48 for adaptation and 37.78 for protected adaptation in normalized units; only the frozen rule crosses the biomass floor. Under substrate-poor startup, protected adaptation yields 86.07 versus frozen 77.13 and truth-informed reference 78.58, demonstrating that the reference is not a finite-horizon optimum. Adaptation instead reduces harvest under optical drift. The lower growth value used for protection is a heuristic, not a validated confidence bound or safety certificate. All figures and performance quantities are reproducible synthetic outputs, with a generated laboratory scene excluded from quantitative evidence.

**Keywords:** continuous fermentation; online adaptation; Monod kinetics; biomass protection; dilution control; synthetic bioprocess

## 1 Continuous operation and the value of adaptation

Continuous fermentation creates a direct coupling between production and biomass retention. Increasing dilution can harvest more cells per unit time, but it also removes the population responsible for converting substrate. A dilution chosen from a nominal steady-state model may remain attractive during stable operation and become excessive after growth activity declines. Online adaptation offers a response, yet an estimator can also react to measurement drift or a transient that its simplified update does not describe well. The question is therefore not whether a parameter changes, but whether changing the dilution on the basis of that estimate improves the complete trajectory. The early chemostat description [1] and continuous-culture analysis [2] establish the physical setting for this balance.

This numerical study asks when a simple online growth update adds value to a steady-productivity policy and whether an additional biomass-dependent dilution cap changes that value. We execute a two-state Monod chemostat with four policies under nominal operation, an activity drop, an affinity change, optical drift and substrate-poor startup. Every quantitative observation, kinetic training pair and process trajectory is synthetic. There is no organism-specific recipe, laboratory culture, pilot vessel or physical sensor run. The model is a controlled example that permits the policy mechanisms to be inspected without attributing performance to an unreported fermentation experiment.

The distinction between steady optimization and transient control is central. A dilution that maximizes equilibrium biomass harvest need not maximize harvest over an 80-hour episode beginning far from equilibrium. An adaptive policy can also use its estimate in a way that is beneficial for one disturbance and harmful for another. Our four rules are deliberately simple: a frozen equilibrium optimum, an online adaptive equilibrium optimum, that adaptive rule with a protective cap and a reference using true kinetic parameters in the same protected construction. None solves a receding-horizon optimization problem. The truth-informed rule is consequently a reference for information access, not an upper bound on achievable finite-horizon production.

The topic connects to a substantial modeling tradition. Hybrid first-principles neural models [3] and process modeling with prior knowledge [4] combine mechanistic structure with data-based components. Recent work uses historical fermentation data for hybrid modeling and predictive control [5], and a 2025 preprint develops adaptive predictive control for fed-batch cultivation [6]. Those methods motivate careful separation of learning, model structure and executed command. The present calculation does not implement their architectures or claim their guarantees. It tests a small transparent adaptation rule that can be reproduced from the equations and code supplied here.

The contribution is an executed comparison rather than a prospective protocol. Twenty-four independent nonlinear kinetic fits initialize 480 control episodes across five conditions and four policies. We report cumulative biomass harvest, substrate outflow and actual biomass-floor crossings. Results include favorable and unfavorable adaptation outcomes and a startup case in which protected adaptation exceeds the truth-informed reference on finite-horizon harvest. The analysis explains that result without treating imperfect parameter estimates as intrinsically superior to true kinetics. A generated laboratory scene provides visual context but supplies no training data, process measurements or performance evidence.

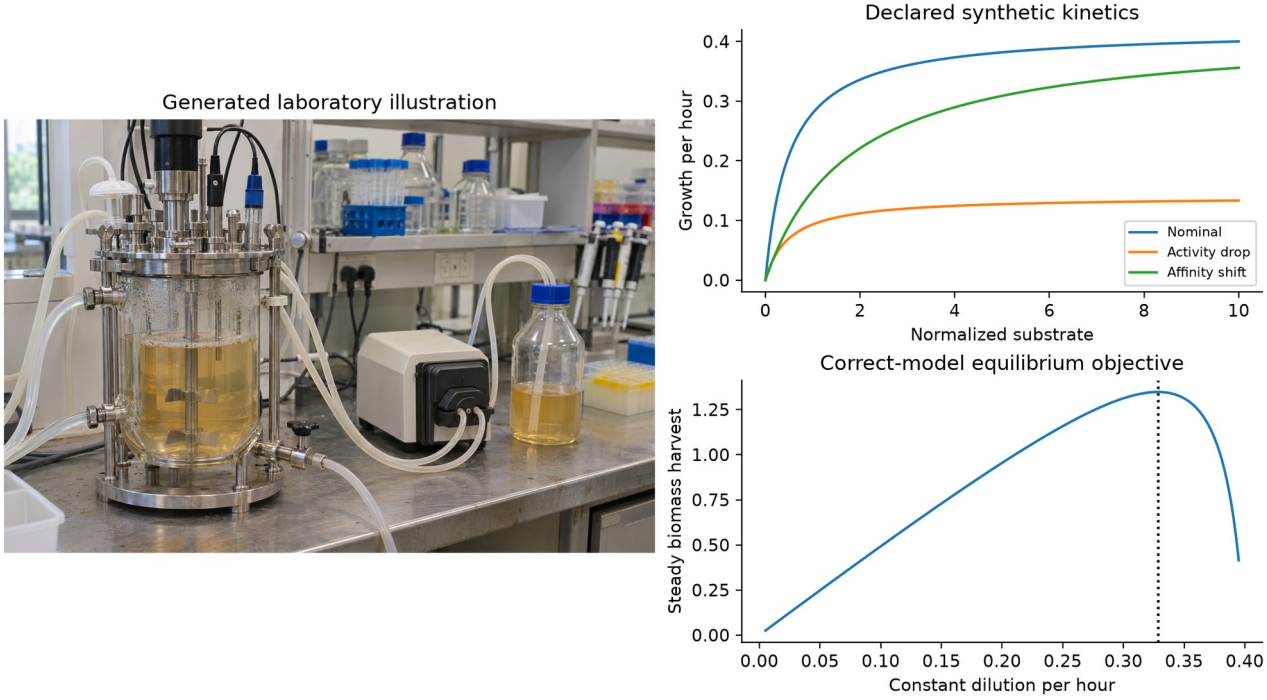



Figure 1. Generated laboratory illustration at left, excluded from all quantitative evidence; executed synthetic kinetic families and equilibrium productivity at right. The photograph-like scene does not depict a performed experiment.

## 2 Synthetic process and equilibrium objective

The plant contains biomass concentration X and growth-limiting substrate concentration S in a constant-volume, perfectly mixed vessel. Dilution D represents matched feed and harvest flow divided by vessel volume. Feed substrate concentration is ten normalized units and biomass yield is 0.5. Specific growth follows a saturating Monod form with maximum growth m and affinity constant K. The nominal values are 0.42 per hour and 0.5 respectively. These normalized concentrations define a numerical process family. They are not calibrated to a species, medium or actual reactor. The time unit is hours, making dilution and growth rates directly comparable within the model.

$$\mu(S)=mS/(K+S) \tag{1}$$

$$dX/dt=[\mu(S)-D]X \tag{2}$$

$$dS/dt=D(S_{in}-S)-\mu(S)X/Y \tag{3}$$

The model omits maintenance, inhibition, product formation and population-density effects. Pirt's maintenance analysis [7], the Contois relation between population density and growth [8] and inhibitory-substrate modeling [9] identify mechanisms outside this construction. Luedeking and Piret's fermentation kinetics [10] also distinguish product formation from biomass growth. We therefore call the integrated biomass outflow harvest rather than claiming a product titer or metabolite yield. These omissions are not hidden inside a learned correction. They delimit the numerical question: how a simple growth-adaptive dilution interacts with biomass and substrate dynamics when the declared plant family is known.

For a positive steady state at a constant dilution below the washout boundary, growth equals dilution. Solving the Monod relation gives the equilibrium substrate concentration, and the substrate balance then gives biomass. Multiplying biomass by dilution yields the steady biomass-harvest rate. Differentiating that expression produces the unconstrained interior optimum used by the steady policies. The optimum depends on the maximum growth and affinity estimates but not on a predicted future trajectory. It is a static operating target. The derivation applies to the declared model and should not be transferred to a fermentation with inhibition, varying yield or additional oxygen constraints without rederiving its objective.

$$S^{*}=KD/(m-D) \tag{4}$$

$$X^{*}=Y(S_{in}-S^{*}) \tag{5}$$

$$D^{*}=m[1-\sqrt{\{K/(K+S_{in})\}}] \tag{6}$$

The nominal optimum is approximately 0.328 per hour. It is feasible below the nominal washout threshold, but it exceeds the maximum growth after the activity-drop condition changes m to 0.14. A frozen policy can then remove biomass faster than even substrate-saturated growth can replace it. This relationship predicts a qualitative failure mechanism before simulation. It does not quantify the finite-horizon loss, the recovery time of an adaptive rule or the effect of a biomass-dependent cap. Those quantities require the executed trajectories. The equilibrium curve in Figure 1 is therefore analytical context, whereas the later figures and tables describe actual numerical episodes.

State integration uses a fourth-order Runge–Kutta scheme with 0.1-hour steps. Each dilution command is held for one hour, during which ten integration steps are taken. Positivity is checked after every step and all executed states remain positive. This is a numerical consistency check, not a proof of positivity for every possible parameter or command. There is no stochastic process noise in the plant. Variability comes from independently fitted training data and noisy biomass observations used by the controller. The deterministic

truth-informed trajectories are consequently identical across replicates in conditions where its state-dependent protection is inactive and observations do not affect commands.

The normal initial state is biomass 0.7 and substrate six. Substrate-poor startup instead begins at biomass 0.48 and substrate 0.02, close to the diagnostic biomass floor 0.45. Initial biomass is above the floor in every case. The floor is not a washout definition or a biological viability threshold. It is an explicit control diagnostic chosen for this synthetic process. A trajectory can remain above it while suffering poor harvest, and crossing it does not immediately make biomass mathematically zero. Keeping this distinction permits the experiment to report production and retention separately rather than equating any one threshold with the whole process outcome.

Table 1. Executed synthetic chemostat and learning design. Biomass and substrate are normalized concentrations.

| Quantity | Value | Role |
|---|---|---|
| Nominal kinetics | Maximum growth 0.42 per h K 0.5 | Monod function with yield 0.5 |
| Feed substrate | 10 normalized units | Constant feed concentration |
| Training per replicate | 100 substrate growth pairs | Noisy kinetic data independent of control |
| Training noise SD | 0.008 per h | Nonlinear least-squares fit of two parameters |
| Control episodes | 24 replicates 5 cases 4 rules | 480 paired numerical policy episodes |
| Integration and control | 0.1 h and 1 h | RK4 state integration and held commands |
| Biomass floor | 0.45 normalized units | Diagnostic threshold used by protected policies |

## 3 Kinetic fitting and online information

Each of 24 replicates begins with 100 independently sampled substrate values between one and nine. Their growth labels follow the nominal Monod relation plus independent normal noise with standard deviation 0.008 per hour. A bounded nonlinear least-squares fit estimates maximum growth and affinity. The lower bounds are 0.05 and 0.03, and upper bounds 0.8 and four respectively. This is an actual executed fit, not a coefficient selected to match a desired control result. Its data are independent of the subsequent control trajectories. The fitted maximum growth has mean approximately 0.42001 and standard deviation 0.00188; the fitted affinity has mean 0.50015 and standard deviation 0.02109.

$$(\hat{m}_0, \hat{K}) = arg\,min_{m,K} \sum_i [g_i - m S_i/(K + S_i)]^2 \tag{7}$$

The substrate sensor is exact in this study. Biomass observation is the actual state multiplied by the exponential of a normal innovation with standard deviation 0.015. The same innovation sequence is shared across the four policies within one replicate and condition, although their states differ because their commands differ. Under optical drift, a deterministic log bias grows by 0.035 per hour after hour 30. The resulting observation remains positive, permitting logarithmic updates. This measurement model does not reproduce an actual optical probe or turbidity calibration. It supplies a known disturbance to distinguish changes in observed biomass from changes in generating growth kinetics.

After each hour, an approximate maximum-growth observation is formed from the change in log observed biomass, the previous held dilution and previous substrate. The update treats the previous substrate as representative of the growth over that interval. That approximation can be poor during a fast substrate transient. Dividing by a small previous substrate amplifies the resulting mismatch as well as observation noise. The code does not conceal this issue by using the true instantaneous growth. It is particularly relevant to the startup experiment, where the substrate begins at 0.02 and changes rapidly after feed enters.

$$\tilde{m}_t = [\log(x_t / x_{t-1}) + D_{t-1}](\hat{K} + S_{t-1}) / S_{t-1} \tag{8}$$

The online maximum-growth estimate is an exponentially weighted update with weights 0.8 and 0.2, clipped between 0.03 and 0.8. The affinity estimate remains frozen at its training value. Thus the procedure has only one adapting kinetic parameter. An affinity change can be absorbed partly into an effective maximum-growth estimate without correctly identifying the underlying kinetic cause. This is a model-misspecification issue in the sense that a fitted parameter can describe a changed relationship without retaining its original physical meaning [11]. It is not evidence that the algorithm has separately estimated both maximum growth and affinity from the episode.

$$\hat{m}_t = clip(0.8\hat{m}_{t-1} + 0.2\tilde{m}_t, 0.03, 0.8) \tag{9}$$

A second exponentially weighted quantity tracks squared update innovations, using weights 0.85 and 0.15. The initial value is the square of 0.008. The heuristic lower growth value subtracts 1.64 times its square root from the updated maximum-growth estimate, with a floor of 0.03. The familiar-looking multiplier does not make this a validated confidence interval. Residuals are not independent parameter-estimation errors, transient approximation bias is present and clipping alters the distribution. We report the fraction of control times at which this lower value lies below true maximum growth only as an empirical diagnostic. It is not a probability guarantee for a future episode.

$$m_{low} = max(0.03, \hat{m}_t - 1.64\sqrt{\{v_t\}}) \tag{10}$$

The online rule is also distinct from a Kalman filter [12]. There is no state covariance propagation, measurement update derived from a likelihood or joint estimation of latent biomass. More elaborate scientific learning methods, including physics-informed networks [13], sparse equation discovery [14] and coordinate discovery [15], address other combinations of structure and data. Citing these approaches provides methodological context, not a claim that they were executed. The present model learning consists of the initial two-parameter nonlinear fit and the explicitly stated scalar online update. Its simplicity makes the information pathway and its failure modes reviewable.

## 4 Dilution policies and protection mechanism

The frozen steady rule applies the equilibrium-optimal dilution computed from its initial fitted parameters throughout the episode. It continues to calculate the online estimate for logging but does not use that estimate in its commands. The adaptive steady rule substitutes the current maximum-growth estimate into the same equilibrium formula while retaining the fitted affinity. Both commands are clipped to the actuator interval from 0.01 to 0.38 per hour. The distinction between calculating an estimate and granting it command authority is important: a diagnostic estimate can be wrong without altering the frozen policy, whereas the same estimate directly changes the adaptive policy's dilution.

Protected adaptation begins with the adaptive steady target and caps it using a state-dependent expression. The cap combines the heuristic lower growth at current substrate with a recovery term that becomes small near the observed biomass floor. If the observed biomass is high, the cap can permit a dilution above the

lower growth value because the recovery term allows some decline. Near the floor it becomes more restrictive. This construction is motivated by the derivative of log biomass relative to a floor, but the implementation uses noisy observations, a heuristic rate and hourly held commands. It is therefore evaluated empirically rather than presented as a formal control-barrier certificate.

$$D \le m_{low} S/(\hat{K}+S) + 0.12(1 - X_{min}/x) \tag{11}$$

After taking the smaller of the steady target and this cap, dilution is again clipped to the actuator interval. The positive lower actuator bound can override a proposed cap below 0.01. This means the mathematical cap is not always enforceable. Moreover, a cap checked at the beginning of one hour does not automatically guarantee its validity throughout that hour. These implementation details matter for interpreting zero observed floor crossings. Formal learning-based control research supplies stronger assumptions and guarantees in other settings [16] [17]. We do not import those guarantees into this heuristic controller merely because it also restricts an adaptive command.

The truth-protected reference uses generating maximum growth and affinity in both the equilibrium target and growth part of the cap. It still uses the same noisy observed biomass in the recovery term. It is therefore truth-informed about kinetics rather than omniscient about every state and future. More importantly, it uses the same steady-target architecture instead of optimizing cumulative harvest over 80 hours. A different transient policy can outperform it on that finite-horizon objective. This is not a paradox or evidence that true parameters are disadvantageous. It distinguishes information quality from the quality of the decision rule that consumes the information.

Five conditions are executed. Nominal kinetics and observation behavior remain unchanged in the first. The activity-drop condition changes maximum growth from 0.42 to 0.14 at hour 30. The affinity-shift condition changes K from 0.5 to 1.8 at that hour. Optical drift leaves plant kinetics nominal and changes the biomass observation bias. Poor startup changes only the initial state. Each condition is evaluated with all four policies for 80 hours in every replicate, giving 480 episodes. The design does not combine disturbances within one episode or search their severities after observing which policy wins.

Cumulative biomass harvest integrates dilution times actual biomass, while substrate outflow integrates dilution times actual substrate. Both are accumulated with trapezoidal integration over the 0.1-hour plant steps. The second is called an outflow proxy rather than waste cost because no price, environmental impact or disposal process is modeled. Floor diagnostics use actual integrated biomass, not its biased observation. An hour is counted if any of its integration steps has biomass below 0.45. The episode crossing fraction is the proportion of 24 replicates with at least one such hour. This definition differs from measuring the exact continuous time spent below the floor.

$$H = \int_0^T D\,X\,dt,\; W = \int_0^T D\,S\,dt \tag{12}$$

## 5 Executed production and retention outcomes

Under nominal operation the policies perform similarly. Mean harvest is 86.335 for the frozen rule, 86.368 for both adaptive rules and 86.341 for the truth-informed reference. No episode crosses the biomass floor. The equality between adaptive and protected adaptation means the cap does not alter their realized commands in this condition. The small harvest difference from the frozen rule is not a convincing universal benefit of adaptation; it is a finite simulated difference under noisy observations and closely fitted nominal parameters. The important finding is that the additional protection does not impose a detectable production penalty here because its constraint is inactive.

After the activity drop, the frozen policy harvest falls to 25.537 and every episode crosses the floor. Its mean minimum biomass is approximately 0.000207, while its fixed dilution remains above the new maximum growth. Adaptive steady operation increases harvest to 34.477 and has no floor crossing in these runs. Protected adaptation increases harvest further to 37.782, with the truth-informed reference reaching 41.422. These values establish that adaptation responds usefully to a genuine kinetic loss under the chosen timing and severity. They also show that protecting the transient can improve total production even when unprotected adaptation already avoids the chosen floor.

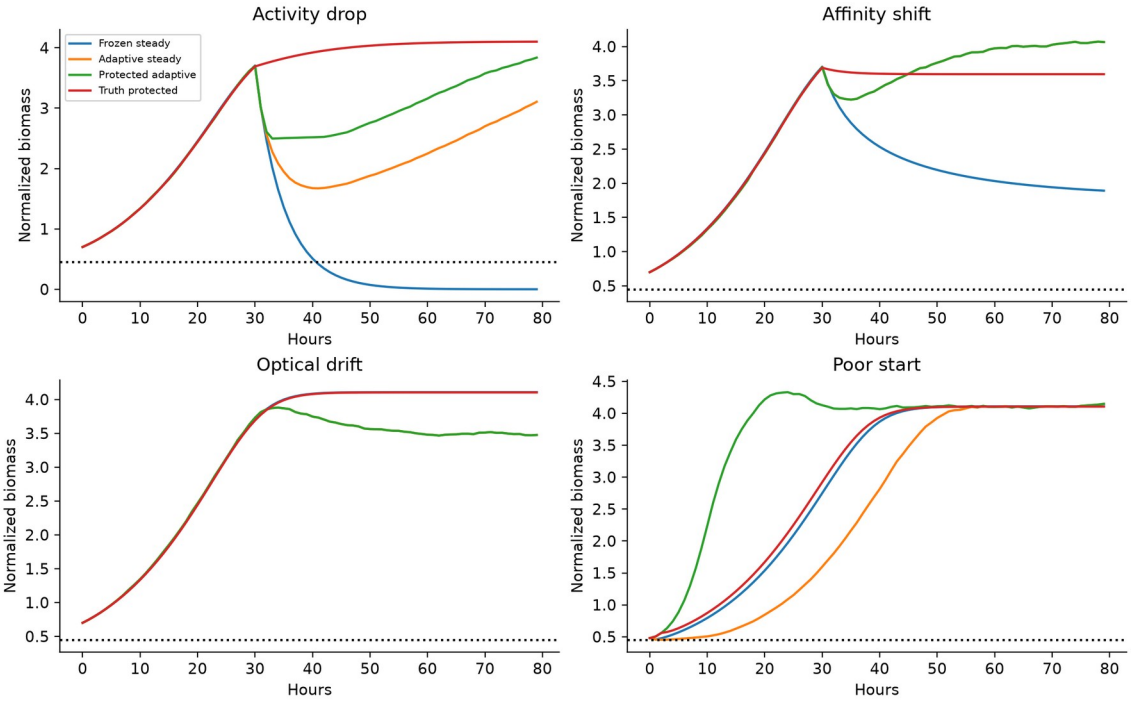



Figure 2. Biomass time courses for the first replicate under four nonnominal conditions. The dotted line is the normalized biomass floor 0.45. Truth-informed protection is a reference policy, not an optimal-control bound.

The activity-drop substrate outflow is 209.011 for the frozen rule, compared with 92.520 for adaptive steady, 81.868 for protected adaptation and 67.785 for the truth-informed reference. The frozen rule continues feeding and discharging substrate as the converting biomass declines. Reducing dilution gives the remaining population time to use substrate and retain biomass. This explanation follows the plant balances and the actual command architecture. It is not a claim of measured resource savings in a factory. The input and output quantities are normalized, and total feed also changes with dilution, so the outflow comparison should not be relabeled as substrate conversion efficiency without accounting for the full material balance.

Table 2. Mean cumulative biomass harvest over 80 hours in normalized units. Columns use the same independently fitted training models and paired observation innovations.

| Condition | Frozen | Adaptive | Protected | Truth reference |
|---|---|---|---|---|
| Nominal | 86.335 | 86.368 | 86.368 | 86.341 |
| Activity drop | 25.537 | 34.477 | 37.782 | 41.422 |
| Affinity shift | 56.414 | 63.971 | 63.971 | 65.584 |

| Condition | Frozen | Adaptive | Protected | Truth reference |
|---|---|---|---|---|
| Optical drift | 86.335 | 83.719 | 83.719 | 86.341 |
| Poor start | 77.125 | 66.460 | 86.073 | 78.583 |

The affinity shift produces a different result. Harvest is 56.414 for frozen operation, 63.971 for either adaptive rule and 65.584 for the truth-informed reference. No policy crosses the floor. Adaptation helps even though it does not update affinity directly. Its effective maximum-growth estimate responds to slower observed growth and changes the steady dilution target. Protected adaptation is identical to adaptive steady because the cap does not become the active command limitation. The result supports useful operational compensation under this specific mismatch. It does not show correct identification of the changed K or justify interpreting the scalar online estimate as a recovered biological maximum.

Optical drift reverses the adaptation benefit. Plant kinetics stay nominal, so the frozen rule retains harvest 86.335. Both adaptive policies instead yield 83.719, while the truth-informed reference gives 86.341. Their substrate outflow rises to 107.036 compared with the frozen 87.414. The observed biomass increase is partly instrumental, but the logarithmic update treats its trend as evidence of growth. In this case the protective cap does not alter the adaptive commands, so it provides no remedy. Zero floor crossings therefore coexist with a production loss. Biomass protection and correct interpretation of the sensor are distinct requirements.

Table 3. Mean cumulative substrate-outflow proxy over 80 hours. This is unconsumed substrate leaving the vessel, not a monetary or environmental impact.

| Condition | Frozen | Adaptive | Protected | Truth reference |
|---|---|---|---|---|
| Nominal | 87.414 | 87.063 | 87.063 | 87.397 |
| Activity drop | 209.011 | 92.520 | 81.868 | 67.785 |
| Affinity shift | 147.257 | 86.810 | 86.810 | 92.718 |
| Optical drift | 87.414 | 107.036 | 107.036 | 87.397 |
| Poor start | 99.413 | 127.091 | 35.661 | 91.241 |

Poor startup produces the most instructive policy ranking. Frozen harvest is 77.125 and adaptive steady harvest falls to 66.460. Protected adaptation reaches 86.073, exceeding the truth-informed reference at 78.583. Substrate outflow is 35.661 for protected adaptation, compared with 99.413, 127.091 and 91.241 for frozen, adaptive and reference rules. No episode crosses the floor, although unprotected minimum biomass approaches it. The initial low substrate makes the scalar update transiently unreliable, and protection retains biomass while limiting early dilution. The observed advantage is a finite-horizon trajectory effect, not a theorem that a more conservative estimate or an imperfect model always produces more harvest.

## 6 Explaining the transient tradeoffs

The startup command trace in Figure 5 makes the mechanism visible. In the first replicate, frozen and adaptive rules initially choose dilution about 0.328 per hour. After the first interval, the approximate maximum-growth observation drives the online estimate to its upper clip 0.8. The adaptive command reaches the actuator maximum 0.38. Protected adaptation instead starts near 0.023 and remains near 0.020 in the next

interval before rising gradually. The truth-informed protected rule also starts near 0.023 but increases more rapidly once substrate rises. These are executed commands from a fixed replicate, not an optimized explanation fitted after the fact.

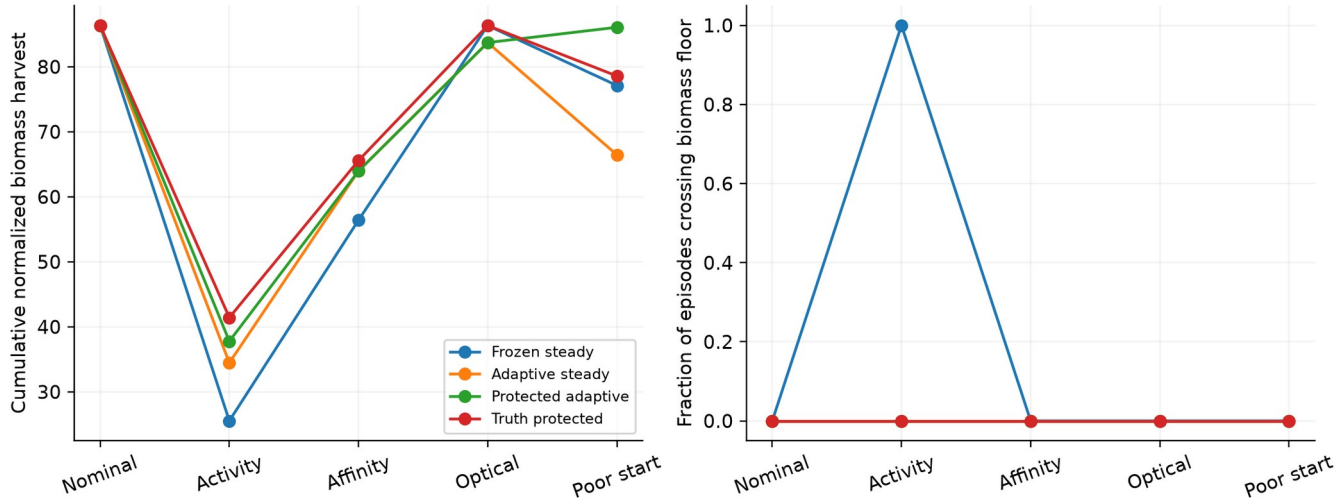



Figure 3. Mean cumulative harvest and fraction of episodes crossing the biomass floor across 24 replicates per condition and policy. Every episode lasts 80 hours.

The startup update fails to describe a fast-changing interval because it divides an average log-growth signal by the very small previous substrate. Substrate increases during that hour, so actual integrated growth is not well represented by the initial substrate value. The estimated maximum growth can therefore be high even without a real kinetic increase. The innovation variance also becomes large, driving the heuristic lower value to its floor. In protected adaptation that conservative quantity restrains dilution for several early hours. The result is beneficial in this particular startup, but the benefit comes from an interaction of estimator approximation and policy structure, not from accurate early parameter recovery.

At the end of poor-start episodes, mean biomass is approximately 4.108 for protected adaptation and 4.104 for the truth-informed reference; final substrate is approximately 1.785 and 1.791. Thus the higher protected harvest is not explained by leaving a dramatically depleted terminal biomass state. The trajectories approach similar final operating regions but spend the startup differently. This check does not turn the heuristic into a finite-horizon optimum. Other policies, horizons or terminal costs could change the ranking. It does show that the reported harvest comparison is more than a final-state accounting artifact within the executed 80-hour episodes.

Figure 4 compares the protected policy's online estimate, heuristic lower value and generating maximum growth during activity loss and optical drift. The lower-value diagnostic lies below the true rate at about 97.5 percent of control times after pooling the complete activity-drop episodes, and about 84.8 percent in optical drift. Those fractions include long periods before or after transitions and are not calibrated coverage probabilities. In the activity-drop case, the truth changes abruptly and the update reacts with delay. In optical drift, the estimator sees a persistent extra log-growth term. A quantity called lower cannot be assumed conservative simply from its algebraic form.

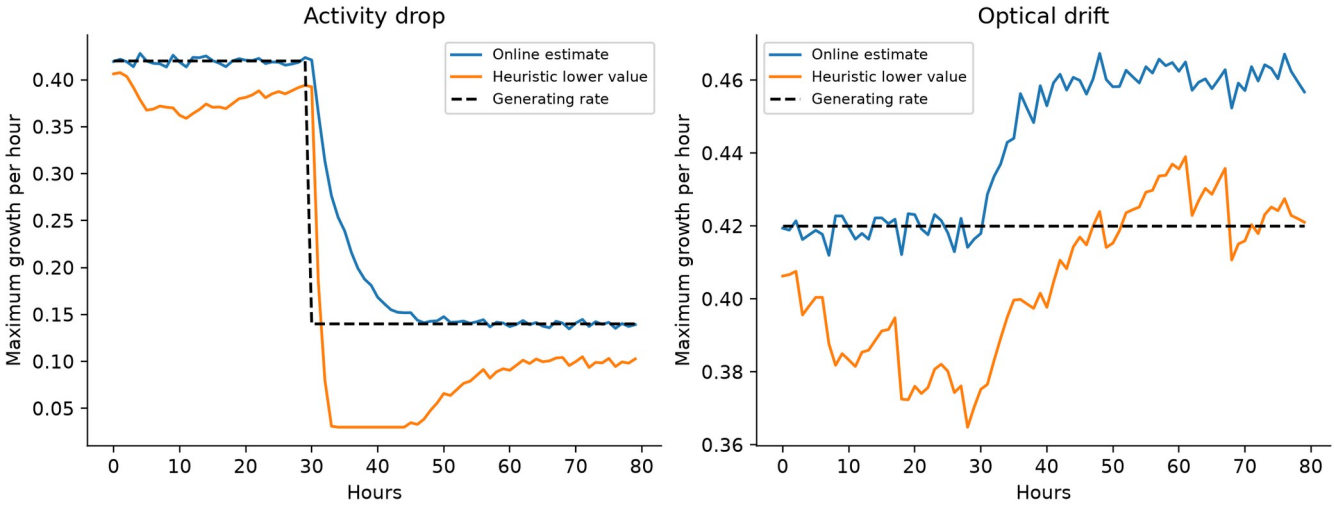



Figure 4. Online maximum-growth estimate, heuristic lower value and generating rate in the protected adaptive policy for activity loss and optical drift. The lower value is not a calibrated statistical confidence bound.

Table 4 records actual biomass outcomes. Only frozen activity-drop episodes cross the floor in this design, with an average of 40 flagged hours. The startup cases remain above it for all policies, but their mean minima differ: approximately 0.4564 for frozen, 0.4534 for adaptive steady and 0.4800 for both protected policies. Reporting only crossing fractions would conceal those margins. Conversely, reporting only the minimum would miss the production cost and substrate behavior. The study therefore retains all three dimensions rather than reducing performance to one safety-versus-productivity slogan. The selected floor is a diagnostic, not a complete definition of acceptable fermentation operation.

Table 4. Biomass outcomes across 24 episodes per row. Minimum is the mean of within-episode minima; crossings and below-floor hours use actual integrated states.

| Condition | Policy | Mean minimum | Crossing fraction | Below-floor hours |
|---|---|---|---|---|
| Nominal | Frozen steady | 0.7000 | 0.000 | 0.00 |
| Nominal | Adaptive steady | 0.7000 | 0.000 | 0.00 |
| Nominal | Protected adaptive | 0.7000 | 0.000 | 0.00 |
| Nominal | Truth protected | 0.7000 | 0.000 | 0.00 |
| Activity drop | Frozen steady | 0.0002 | 1.000 | 40.00 |
| Activity drop | Adaptive steady | 0.7000 | 0.000 | 0.00 |
| Activity drop | Protected adaptive | 0.7000 | 0.000 | 0.00 |
| Activity drop | Truth protected | 0.7000 | 0.000 | 0.00 |

| Condition | Policy | Mean minimum | Crossing fraction | Below-floor hours |
|---|---|---|---|---|
| Affinity shift | Frozen steady | 0.7000 | 0.000 | 0.00 |
| Affinity shift | Adaptive steady | 0.7000 | 0.000 | 0.00 |
| Affinity shift | Protected adaptive | 0.7000 | 0.000 | 0.00 |
| Affinity shift | Truth protected | 0.7000 | 0.000 | 0.00 |
| Optical drift | Frozen steady | 0.7000 | 0.000 | 0.00 |
| Optical drift | Adaptive steady | 0.7000 | 0.000 | 0.00 |
| Optical drift | Protected adaptive | 0.7000 | 0.000 | 0.00 |
| Optical drift | Truth protected | 0.7000 | 0.000 | 0.00 |
| Poor start | Frozen steady | 0.4564 | 0.000 | 0.00 |
| Poor start | Adaptive steady | 0.4534 | 0.000 | 0.00 |
| Poor start | Protected adaptive | 0.4800 | 0.000 | 0.00 |
| Poor start | Truth protected | 0.4800 | 0.000 | 0.00 |

## 7 Relation to bioprocess learning and control

Recent primary work demonstrates several more elaborate control directions. Batch-to-batch optimization using recursively updated extreme learning machines [18] and a fermentation digital-twin case study [19] address learning at different time scales and with different process representations. Historical-data hybrid modeling for recombinant fermentation [5] connects data availability to predictive control development. The adaptive fed-batch preprint [6] explicitly concerns a different process configuration and controller. These papers are cited within the 2025 horizon as domain research, while the present chemostat uses only the disclosed steady rule and scalar adaptation. It would be inaccurate to describe its numerical results as evidence for a neural predictive controller.

Primary 2024 studies also span temperature and feed-related control. Ethanol temperature control with a particle-swarm-tuned PID strategy [20], lysozyme temperature control [21] and reinforcement-learning bioreactor temperature control [22] concern manipulated variables and dynamics absent from our two-state model. Koopman modeling with switching operators [23] and predictive control with cell recycling [24] use additional structural information. We do not transfer their empirical or simulation advantages to our dilution policies. Their relevance is that the meaning of a learned model depends on the process configuration, command authority and performance quantity that is actually evaluated.

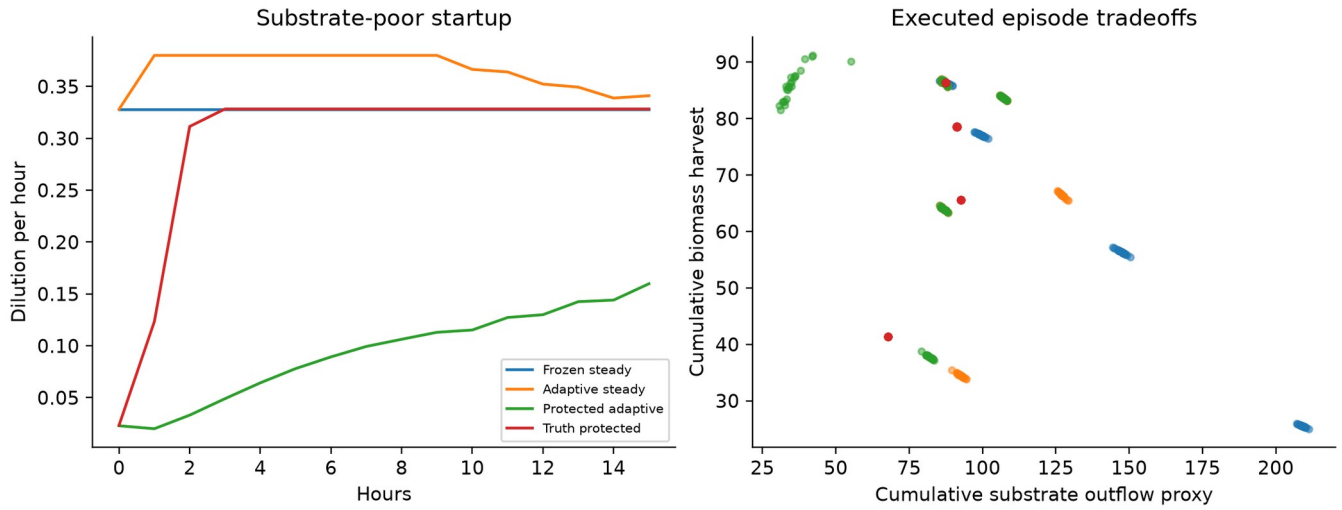



Figure 5. Early dilution commands during substrate-poor startup and episode-level cumulative substrate-outflow versus biomass-harvest tradeoffs. All axes use normalized synthetic quantities except time and dilution rate.

A globally stabilizing continuous-ethanol controller [25] makes the startup issue particularly salient, while fermentation-kinetics optimization for lysine [26] illustrates the coupling of fitted kinetics and operating decisions. Biosensing microchips for ethanol inhibition control [27] concern a physical measurement technology rather than the abstract optical-drift model used here. These studies are not substitutes for validating our heuristic, and our simulated sensor is not a model of those devices. The comparison instead shows why learning and sensing must be evaluated through their effects on a trajectory. A parameter estimate can look responsive while directing the plant toward a less useful operating history.

The distinction between protection and optimization should remain explicit. The biomass cap is a local rule motivated by a derivative condition, while the objective used to propose dilution is steady harvest. Formal safe learning and data-driven predictive-control results [16] [17] require specific uncertainty, feasibility and sampling assumptions. Our noise, approximate update, actuator floor and hourly hold do not satisfy a demonstrated theorem of that kind. Zero observed crossings are finite experimental outcomes. They support comparing the executed policies, not claiming certified safety, robust invariance or guaranteed recovery for every initial state and disturbance.

Computationally, the experiment uses a standard nonlinear least-squares implementation and explicit integration, within the scientific-computing ecosystem described by SciPy [28]. The code records the initial fits, all episode metrics and complete first-replicate paths. There is no hidden training schedule or policy search. Parameters controlling update weights, clipping and protection are fixed before the factorial episodes. A different tuning could change outcomes, but we do not choose it using the test results. This simple structure permits the observed optical-drift loss and startup gain to remain part of the same evidence rather than being removed through condition-specific retuning.

## 8 Parameter information and finite horizon decisions

The initial kinetic fit is deliberately favorable but not noiseless. Training substrates lie between one and nine, mostly above the nominal affinity value 0.5. In that range the growth curve approaches its maximum, so maximum growth can be estimated more tightly than affinity. The fitted standard deviations reflect this difference: approximately 0.00188 for maximum growth and 0.02109 for affinity. These are across-replicate fitting variations, not posterior uncertainty intervals supplied to the controller. The online protection variance

begins from a separate fixed number and then follows innovations. It should not be interpreted as propagating the covariance of the original two-parameter nonlinear fit.

The affinity-shift experiment highlights practical compensation without mechanistic identification. When K rises, the same substrate supports less specific growth. The scalar update still divides by the old fitted K relationship and returns an effective maximum-growth value. Lowering dilution in response can improve harvest even though the estimate has the wrong physical interpretation. This is useful for control within a limited operating range, but it can fail when the process moves to a different substrate regime. A single effective parameter cannot generally reproduce both the low-substrate slope and high-substrate plateau of a changed Monod curve. The reported benefit therefore concerns executed operational adaptation, not discovery of the true changed kinetics.

The activity-loss case is different because its maximum-growth change lies within the adapting parameter's meaning. The update has a more direct path to compensating the disturbance. Even here it reacts through noisy finite differences and a smoothing lag, so protection can change the transient while the estimate catches up. Relative to the frozen rule, protected adaptation raises mean harvest by approximately 47.95 percent in this condition and wins in all 24 paired replicates. Under affinity shift the corresponding gain is approximately 13.40 percent, again with 24 paired wins. These counts describe the executed disturbance family and paired innovations. They do not establish an independent probability of success for a real fermentation campaign.

Optical drift gives 24 paired losses for protected adaptation relative to frozen operation, with mean harvest reduction approximately 3.03 percent. This result remains in the same comparison because no condition-specific tuning is allowed. A controller that adapts only when a genuine biological change occurs would need information that the current scalar observation pathway does not possess. The simulation supplies exact substrate but no independent biomass assay or sensor-drift estimator. A persistent rise in log observed biomass can therefore be interpreted as growth. The biomass cap is inactive and cannot solve this identification ambiguity. Protecting a state threshold and validating a measurement trend are separate mechanisms in the executed architecture.

Protected adaptation wins all 24 poor-start comparisons with frozen operation, with mean harvest gain approximately 11.60 percent. Its harvest standard deviation is 2.8530, much larger than the 0.2880 frozen-rule value. The advantage is therefore accompanied by greater sensitivity to the sampled training fit and observation sequence. Reporting only the mean would conceal that spread. The supplied example path helps explain one realization, while the replicate outputs retain the distribution. We do not claim that the protection rule is tuned to minimize variance or that its startup result will persist for every initial substrate, biomass floor or adaptation gain.

A material-balance interpretation prevents another overstatement. In this constant-volume model the combination of substrate and biomass expressed in substrate-equivalent units changes through feed and outflow. Increasing dilution changes the amount of feed entering over 80 hours as well as the harvested biomass and unconsumed substrate leaving. Cumulative harvest divided by substrate outflow is consequently not a conversion yield. The latter would require a denominator accounting for feed, initial inventory and terminal inventory. We report harvest and outflow as separate observed simulation integrals. The terminal-state check supports comparison of trajectories, but it does not replace a complete resource or economic accounting model.

The protective cap also changes substrate residence time. Lower early dilution retains biomass but admits feed more slowly. Depending on the initial state, either effect could dominate. In the executed poor start, the

cap allows biomass to increase while substrate remains relatively low, then releases dilution gradually. The truth-informed steady target is released sooner and moves the vessel through a different startup history. That difference explains how a policy with poorer early parameter estimates can accumulate more harvest over a fixed horizon. It is the decision architecture and transient path that matter. Replacing the reference with an actual finite-horizon optimizer would be a different comparison and could produce a different ordering.

The hourly command period is part of this result. The log update uses a full interval, while the plant can change substantially inside it. Faster sampling could reduce some transient approximation error, but it could also change the noise amplification in finite differences. Slower sampling would produce another balance between integrated growth and the previous substrate approximation. We do not vary the command period and therefore do not infer its optimal value. The declared one-hour hold and 0.1-hour integration step specify the actual numerical system whose performance is reported. This precision matters because a continuous-time protective inequality is only one ingredient of a sampled implementation.

The actuator clipping is equally substantive. An upper limit of 0.38 prevents the adaptive rule from executing arbitrarily large dilution after an inflated estimate. A lower limit of 0.01 means the system never shuts feed and harvest completely. The lower clip can override a very restrictive cap, while the upper clip can limit a nominal steady target. Both are applied to the final command rather than merely described as desired operating limits. No rate-of-change constraint or pump dynamics is modeled. The trajectories therefore include exact held dilution changes that a physical actuator might not reproduce. This is another boundary between the synthetic mechanism and a deployable fermentation controller.

The five conditions also should not be pooled into one unqualified controller ranking. Their frequencies are experimental design choices, not a distribution of expected plant incidents. Averaging all episodes would assign equal weight to nominal operation, a severe activity loss, an affinity change, optical drift and one startup. A different mix could favor a different policy. Reporting the factorial rows preserves the conditions under which adaptation helps or hurts. The episode tradeoff scatter is descriptive of those 480 numerical runs, not an estimate of the long-run operating population. Choosing a preferred controller would require the intended operating scenarios and objectives to be stated separately.

Finally, the study illustrates how to make a compact adaptive-control claim reviewable. The learned quantities, observation model, command formula, integration scheme and actual evaluation target are all explicit. The same estimator can be logged without command authority, used in a steady target or restricted by another rule, enabling a direct comparison of information and action. Favorable production, unfavorable drift response and a nonoptimal truth-informed reference remain visible together. That combination is the scientific result. It supports investigating richer estimation or optimization while avoiding the unsupported conclusion that adding learning, adding protection or supplying true parameters must improve every metric in every transient.

The replicate definition is also important. The 24 fits and innovation sequences do not represent different microbial strains or independent laboratory campaigns. They are numerical realizations of one nominal kinetic-training experiment followed by paired policy episodes. True plant parameters are fixed within each declared condition apart from its specified change. The observed spread therefore measures sensitivity to those synthetic data and observation draws, not biological diversity. This restricted interpretation makes the paired comparisons useful without overstating their scope. It also explains why a truth-informed policy can have zero harvest variation when its executed command is independent of those draws.

## 9 Limits and reproducible conclusions

The model excludes many important bioprocess constraints. Dissolved oxygen, pH, temperature, mixing, contamination, product inhibition, byproduct formation and biomass composition are absent. Feed concentration and yield are fixed, and substrate is observed without noise. The plant uses the same growth family as the initial fit, except for declared parameter changes. Real fermentation can violate those assumptions in multiple ways at once. The present results are therefore mechanistic synthetic evidence for a small control architecture. They do not establish a viable operating recipe, a physical sensor specification or expected production improvement for any actual organism.

The protection quantity is heuristic and its lower clipping value can exceed an actual rate in a more severe disturbance. The actuator minimum can also prevent execution of the proposed cap. No explicit feasibility test, emergency shutdown or guaranteed fallback is modeled. These details prevent a formal safety claim even though protected policies avoid the selected floor in all executed conditions. Similarly, the truth-informed reference is neither a perfect-state observer nor a finite-horizon optimizer. Its lower startup harvest is an architectural result. Interpreting it as an upper bound would incorrectly convert a useful comparison into a claim about optimal control that the code does not solve.

Reproduction requires NumPy, pandas, SciPy and Matplotlib. The self-contained script generates 24 kinetic datasets, fits their parameters, executes all 480 episodes and writes training_fits.csv, runs.csv, summary.csv, example_paths.csv and results.json. It also regenerates the five figures, using the retained generated laboratory illustration only in Figure 1. The image provenance contains its exact prompt and the distinction between illustration and numerical evidence. No private process data, external model service or host-specific dependency path is required. The saved records permit verification of both favorable and unfavorable policy outcomes, including the apparently surprising startup comparison.

The conclusion is conditional: online growth adaptation helps when the scalar update tracks a genuine activity change, can compensate operationally for a frozen affinity mismatch and can hurt when an observation drift is interpreted as growth. Biomass-dependent protection changes transient dilution and can improve finite-horizon harvest without being a certified safety mechanism. Its startup advantage over a truth-informed steady reference shows why better kinetic knowledge and better transient decisions are separate questions. The executed study provides a reproducible baseline for examining that distinction before introducing richer models, additional state measurements or predictive optimization.

## References

[1] Novick, Aaron; Szilard, Leo (1950). Description of the Chemostat. Science, 112, 2920, 715-716. DOI: 10.1126/science.112.2920.715.

[2] Herbert, D.; Elsworth, R.; Telling, R. C. (1956). The Continuous Culture of Bacteria; a Theoretical and Experimental Study. Journal of General Microbiology, 14, 3, 601-622. DOI: 10.1099/00221287-14-3-601.

[3] Psichogios, Dimitris C.; Ungar, Lyle H. (1992). A hybrid neural network-first principles approach to process modeling. AIChE Journal, 38, 10, 1499-1511. DOI: 10.1002/aic.690381003.

[4] Thompson, Michael L.; Kramer, Mark A. (1994). Modeling chemical processes using prior knowledge and neural networks. AIChE Journal, 40, 8, 1328-1340. DOI: 10.1002/aic.690400806.

[5] Bolmanis, Emils; Galvanauskas, Vytautas; Grigs, Oskars; Vanags, Juris; Kazaks, Andris (2025). Leveraging Historical Process Data for Recombinant P. pastoris Fermentation Hybrid Deep Modeling and Model Predictive Control Development. Fermentation, 11, 7, 411. DOI: 10.3390/fermentation11070411.

[6] Krausch, Niels; Doff-Sotta, Martin; Cannon, Mark; Neubauer, Peter; Bournazou, Mariano Nicolas Cruz (2025). Deep learning adaptive Model Predictive Control of Fed-Batch Cultivations. arXiv preprint arXiv:2502.01488, submitted 3 February 2025.

[7] Pirt, S. J. (1965). The maintenance energy of bacteria in growing cultures. Proceedings of the Royal Society of London. Series B. Biological Sciences, 163, 991, 224-231. DOI: 10.1098/rspb.1965.0069.

[8] CONTOIS, D. E. (1959). Kinetics of Bacterial Growth: Relationship between Population Density and Specific Growth Rate of Continuous Cultures. Journal of General Microbiology, 21, 1, 40-50. DOI: 10.1099/00221287-21-1-40.

[9] Andrews, John F. (1968). A mathematical model for the continuous culture of microorganisms utilizing inhibitory substrates. Biotechnology and Bioengineering, 10, 6, 707-723. DOI: 10.1002/bit.260100602.

[10] Luedeking, Robert; Piret, Edgar L. (1959). A kinetic study of the lactic acid fermentation. Batch process at controlled pH. Journal of Biochemical and Microbiological Technology and Engineering, 1, 4, 393-412. DOI: 10.1002/jbmte.390010406.

[11] H. White, "Maximum Likelihood Estimation of Misspecified Models," Econometrica, vol. 50, no. 1, pp. 1, 1982. doi: 10.2307/1912526.

[12] R. E. Kalman (1960). A New Approach to Linear Filtering and Prediction Problems. Journal of Basic Engineering. https://doi.org/10.1115/1.3662552

[13] Raissi M., Perdikaris P., Karniadakis G.E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378, 686-707. https://doi.org/10.1016/j.jcp.2018.10.045

[14] Brunton Steven L., Proctor Joshua L., Kutz J. Nathan (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the National Academy of Sciences, 113, 3932-3937. https://doi.org/10.1073/pnas.1517384113

[15] Champion, Kathleen; Lusch, Bethany; Kutz, J. Nathan; Brunton, Steven L. (2019). Data-driven discovery of coordinates and governing equations. Proceedings of the National Academy of Sciences, 116, 45, 22445-22451. DOI: 10.1073/pnas.1906995116.

[16] Aswani, Anil; Gonzalez, Humberto; Sastry, S. Shankar; Tomlin, Claire (2013). Provably safe and robust learning-based model predictive control. Automatica, 49, 5, 1216-1226. DOI: 10.1016/j.automatica.2013.02.003.

[17] Berberich, Julian; Kohler, Johannes; Muller, Matthias A.; Allgower, Frank (2021). Data-Driven Model Predictive Control With Stability and Robustness Guarantees. IEEE Transactions on Automatic Control, 66, 4, 1702-1717. DOI: 10.1109/tac.2020.3000182.

[18] Moore, Alex; Zhang, Jie (2025). Batch-to-Batch Optimization Control of Fed-Batch Fermentation Process Based on Recursively Updated Extreme Learning Machine Models. Algorithms, 18, 2, 87. DOI: 10.3390/a18020087.

[19] Zhao, Songguang; Jiao, Tianhui; Adade, Selorm Yao-Say Solomon; Wang, Zhen; Ouyang, Qin; Chen, Quansheng (2025). Digital twin for predicting and controlling food fermentation: A case study of kombucha fermentation. Journal of Food Engineering, 393, 112467, 112467. DOI: 10.1016/j.jfoodeng.2025.112467.

[20] Yerolla, Raju; P, Suhailam; Besta, Chandra Shekar (2024). Advanced temperature control in ethanol fermentation using a PSO-PID controller with split-range control strategy. Preparative Biochemistry & Biotechnology, 55, 2, 196-209. DOI: 10.1080/10826068.2024.2381761.

[21] Ding, Chenhua; Li, Xungen; Zhou, Hanlin; Yu, Jianming; Du, Juling; Zhao, Shixiang (2024). A Temperature Control Method of Lysozyme Fermentation Based on LRWOA-LSTM-PID. Processes, 12, 5, 866. DOI: 10.3390/pr12050866.

[22] Rajasekhar, Nagabhushanamgari; Radhakrishnan, Thota Karunakaran; Mohamed, Samsudeen Naina (2024). Reinforcement learning based temperature control of a fermentation bioreactor for ethanol production. Biotechnology and Bioengineering, 121, 10, 3114-3127. DOI: 10.1002/bit.28784.

[23] Yuan, Jinlong; Zhao, Shuang; Yang, Dongyao; Liu, Chongyang; Wu, Changzhi; Zhou, Tao; Lin, Sida; Zhang, Yuduo; Cheng, Wanli (2024). Koopman modeling and optimal control for microbial fed-batch fermentation with switching operators. Nonlinear Analysis: Hybrid Systems, 52, 101461, 101461. DOI: 10.1016/j.nahs.2023.101461.

[24] Bannoud, Mohamad Al; Ferreira, Paulo Henrique Nascimento; de Andrade, Rafael Ramos; da Silva, Carlos Alexandre Moreira (2024). Control of an integrated first and second-generation continuous alcoholic fermentation process with cell recycling using model predictive control. Chemical Engineering Communications, 212, 4, 521-544. DOI: 10.1080/00986445.2024.2417901.

[25] Qin, Yuxin; Zhai, Chi (2024). Global Stabilizing Control of a Continuous Ethanol Fermentation Process Starting from Batch Mode Production. Processes, 12, 4, 819. DOI: 10.3390/pr12040819.

[26] Li, Hui; Chen, Jiajun; Li, Xingyan; Gan, Jian; Liu, Huazong; Jian, Zhou; Xu, Sheng; Zhang, Alei; Li, Ganlu; Chen, Kequan (2024). Artificial neural network and genetic algorithm coupled fermentation kinetics to regulate L-lysine fermentation. Bioresource Technology, 393, 130151, 130151. DOI: 10.1016/j.biortech.2023.130151.

[27] Zhang, Shaoqi; Wang, Meiyue; Xie, Ying; Li, Shuhan; Chen, Ying; Wu, Hao; Cheng, Donghao; Chu, Zhenyu; Jin, Wanqin (2024). Ultrasensitive Biosensing Microchips to Control Ethanol Fermentation for Effectively Reducing Product Inhibition. Industrial & Engineering Chemistry Research, 63, 49, 21168-21179. DOI: 10.1021/acs.iecr.4c02595.

[28] Virtanen, Pauli; Gommers, Ralf; Oliphant, Travis E.; Haberland, Matt; Reddy, Tyler; Cournapeau, David; Burovski, Evgeni; Peterson, Pearu; Weckesser, Warren; Bright, Jonathan; van der Walt, Stéfan J.; Brett, Matthew; Wilson, Joshua; Millman, K. Jarrod; Mayorov, Nikolay; Nelson, Andrew R. J.; Jones, Eric; Kern, Robert; Larson, Eric; Carey, C J; Polat, İlhan; Feng, Yu; Moore, Eric W.; VanderPlas, Jake; Laxalde, Denis; Perktold, Josef; Cimrman, Robert; Henriksen, Ian; Quintero, E. A.; Harris, Charles R.; Archibald, Anne M.; Ribeiro, Antônio H.; Pedregosa, Fabian; van Mulbregt, Paul; , ; Vijaykumar, Aditya; Bardelli, Alessandro Pietro; Rothberg, Alex; Hilboll, Andreas; Kloeckner, Andreas; Scopatz, Anthony; Lee, Antony; Rokem, Ariel; Woods, C. Nathan; Fulton, Chad; Masson, Charles; Häggström, Christian; Fitzgerald, Clark; Nicholson, David A.; Hagen, David R.; Pasechnik, Dmitrii V.; Olivetti, Emanuele; Martin, Eric; Wieser, Eric; Silva, Fabrice; Lenders, Felix; Wilhelm, Florian; Young, G.; Price, Gavin A.; Ingold, Gert-Ludwig; Allen, Gregory E.; Lee, Gregory R.; Audren, Hervé; Probst, Irvin; Dietrich, Jörg P.; Silterra, Jacob; Webber, James T; Slavič, Janko; Nothman, Joel; Buchner, Johannes; Kulick, Johannes; Schönberger, Johannes L.; de Miranda Cardoso, José Vinícius; Reimer, Joscha; Harrington, Joseph; Rodríguez, Juan Luis Cano; Nunez-Iglesias, Juan; Kuczynski, Justin; Tritz, Kevin; Thoma, Martin; Newville, Matthew; Kümmerer, Matthias; Bolingbroke, Maximilian; Tartre, Michael; Pak, Mikhail; Smith, Nathaniel J.; Nowaczyk, Nikolai; Shebanov, Nikolay; Pavlyk, Oleksandr; Brodtkorb, Per A.; Lee, Perry; McGibbon, Robert T.; Feldbauer, Roman; Lewis, Sam; Tygier, Sam; Sievert, Scott; Vigna, Sebastiano; Peterson, Stefan; More, Surhud; Pudlik, Tadeusz; Oshima, Takuya; Pingel, Thomas J.; Robitaille, Thomas P.; Spura, Thomas; Jones, Thouis R.; Cera, Tim; Leslie, Tim; Zito, Tiziano; Krauss, Tom; Upadhyay, Utkarsh; Halchenko, Yaroslav O.; Vázquez-Baeza, Yoshiki (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17, 3, 261-272. DOI: 10.1038/s41592-019-0686-2.