International Core Journal of Engineering Volume 12 Issue 11, 2026
ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

Bayesian Optimization for Tuning Nonlinear Predictive Control of an Acetone Purification Column

Yiliang Chen, Mingge Sun*, Xiangyu Zhang

School of Information and Control Engineering, Jilin University of Chemical Technology, Jilin, 132022, China

*Corresponding author: Mingge Sun

Abstract

A Bayesian optimization method is developed to tune predictive control for an acetone purification column operating in a high-purity region. Aspen Dynamics is coupled with MATLAB through OPC DA. With the reboiler duty fixed, reflux flow is manipulated using an incremental ridge regression model identified from 20 historical trajectories in logarithmic impurity coordinates. A nonlinear inverse transformation provides overhead composition predictions. The inner controller optimizes three move blocks subject to magnitude and rate limits, while an outer Gaussian process and expected improvement criterion select the move penalty with post-run purity-margin screening. Fifteen valid co-simulation runs include relative feed-flow change up to ±0.1% and relative feed-acetone mole-fraction decrease up to 0.1%. In the negative-flow confirmation run, the frozen penalty value of 0.01 has a minimum overhead acetone mole fraction value of 0.999004900 and causes a 22.41% reduction in the fixed performance score value relative to the mean of two runs with an original penalty value of 0.05. All valid samples from the three frozen-parameter confirmation runs meet the 0.999 criterion; but too much purity still remains when faced with the positive-flow disturbance. These results indicate sampled product-quality satisfaction and higher aggregate performance in the negative-flow case, within the local disturbance range and operating limits investigated here.

Keywords

Acetone Purification Column; Nonlinear Predictive Control; Bayesian Optimization; Logarithmic Impurity Model; Controller Tuning; Co-Simulation.

1. Introduction

Moreover, IPA dehydrogenation process has the downstream separation requirement to obtain the required product quality acetone. Material and energy transport within a distillation column leads to longer dynamics, and small impurity changes in the high-purity region can substantially alter the available quality margin. Studies of distillation dynamics and control show that control structure, internal flows, and transformations of high-purity composition coordinates affect both dynamic models and controller performance[1]. Maintaining overhead product quality within a prescribed operating envelope therefore requires consideration of prediction-variable scaling, manipulated-variable limits, and disturbance effects.

Model predictive control (MPC) explicitly handles operating constraints by repeatedly solving a finite-horizon optimization problem[2]. Data-based modeling provides control-oriented predictors for complex processes, but a lower prediction error does not automatically imply better closed-loop performance. The prediction model, cost parameters, and constraint design must be considered

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

together. Hewing et al. distinguish model learning, controller-parameter learning, and constraint-oriented approaches within learning-based MPC, and identify controller-parameter learning as a means of directly improving closed-loop performance[3]. Input-output history representations provide a common basis for data-based predictors[4], while regularized regression addresses parameter estimation with correlated features[5]. Accordingly, this study uses a low-complexity incremental model and examines its applicability to a local high-purity operating region through multistep rollouts and closed-loop tests.

Closed-loop weight tuning generally requires complete dynamic runs, making each performance evaluation substantially more expensive than an individual control calculation. Surrogate models and acquisition criteria provide a basis for allocating a limited budget of expensive function evaluations[6-8]. Marco et al. used Gaussian process global optimization for automatic controller tuning, demonstrating the feasibility of improving controller parameters using experimentally evaluated performance[9]. Kudva et al. investigated sparse-subspace surrogates for high-dimensional MPC tuning to address the relationship between sample requirements and search dimensionalit[10]. For processes with unknown constraints, Xu et al. developed contextual Bayesian optimization with a constraint-violation budget[11], whereas Gardner et al. modeled the objective and constraints separately to address inequality constraints[12]. These studies also distinguish manipulated-variable limits, post-run feasibility screening, and safety-aware acquisition based on probabilistic constraint models; these mechanisms are not interchangeable.

This study considers a single overhead-quality objective for purification of acetone in column T2. It considers predictive reflux control with fixed reboiler duty and Bayesian optimization (BO) of one tuning parameter. The work comprises an incremental logarithmic-impurity predictor with a nonlinear output transformation, a discrete three-block receding-horizon controller subject to operating constraints, and Gaussian process (GP) and expected improvement (EI) based tuning under a fixed external performance score. The parameter-selection data are kept separate from runs after parameter selection and fixed for those runs. Scope includes local co-simulation under a finite evaluation budget. Post-run purity screening is viewed as a quality criterion for an engineering acceptance test instead of being a safety certificate, and performance is evaluated separately for each disturbance.

2. Process and Co-simulation Platform

2.1 Quality Objective and Operating Conditions

The process is a purifier column T2 in a IPA dehydrogenation flowsheet. The model of the column contains 67 stages, the feed coming in stage 54. The primary quality variable is the overhead acetone mole fraction, x_D , with a quality product threshold of $x_D \geq 0.999$, the tracking reference $r = 0.999005$. The tracking reference is 5×10^{-6} above product threshold level to provide margin for local perturbation. The recorded bottom, x_B, is the mole fraction of acetone; no bottom-quality constraint is imposed in this work. Stage-57 temperature T57 is used only for checking the initial condition and tracking the process.

Reflux flow L is the only online manipulated variable; reboiler duty QR is fixed at 1177720 W. Measured feed-flow, FF, is considered a measured disturbance, and feed-flow controller setpoint FSP is used to impose flow steps. Table 1 summarizes this configuration. The normalization base, $L_{nom} = 100.637$ kmol/h, which is different from the initial reflux-flow value for these tests. The reflux envelope extends from −0.02% to +0.05% relative to L_{nom}, and the requested change over 20 s is limited to $s_L = 0.0100637$ kmol/h.

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

Table 1. Column and controller configuration

Parameter	Value	Parameter	Value
Stages / feed stage	67 / 54	QR	1177720 W
Nominal FF	76.2270138 kmol/h	Nominal L	100.637 kmol/h
Lower L limit	100.6168726 kmol/h	Upper L limit	100.6873185 kmol/h
Sampling interval	2 s	Control interval	20 s
Prediction horizon	60 s / 30 steps	Control blocks	3 blocks / 20 s each
Maximum requested move	0.0100637 kmol/h	Tracking reference	0.999005
Initial L	100.63396454 kmol/h	Initial xD	Approx. 0.99900538
Initial T57	Approx. 73.06 °C	Original move penalty	0.05

2.2 Communication and Common Dynamic Initial Condition

The co-simulator includes the Aspen Dynamics flowsheet, a KEPServerEX V6 OPC DA server, and the MATLAB R2023b controller running on a Windows 10 machine. The program reads L, QR, xD, xB, T57, FF,FSP and the 4 feed-component fractions and writes the reflux commands or disturbance tags if necessary. Each sample is appended to the data file, together with information on the quality of communications, acquisition wall time and readbacks of commands. Check validity of quality of communication and finiteness of numerical values. If write exception takes place, operations subsequent to the control command are inhibited and stored available records are kept; a control command with inconclusive execution is not retaken.

All the comparisons in closed-loop calculations start by loading the full dynamic snapshot S_margin_candidate. This snapshot was acquired by a disturbance-free run to determine the overhead purity margin. Available read-only checks of the initial L, QR, xD, T57, FF, and four feed-component fractions must pass before a test starts. Even though the overhead composition had become reasonably smooth when the snapshot was stored, steady state of all internal column variables was not reached. The snapshot is therefore considered as a typical dynamic initial condition. Agreement between a checked tag and its tag reference is a necessary startup condition but does not fulfill a restoration of the complete snapshot.

3. Logarithmic Impurity Predictor and Controller

3.1 Incremental Model and Parameter Identification

We use a logarithmic impurity coordinate to reflect changes in small composition in the region of high purity. Equation (1) does this. Here ε = 10−6 so that the transformation is finite as composition approaches any pure component. It is a numerical regularization constant and not an assurance of measurement precision. Twenty historical samples of each variable are used. Histories of q, L, QR, and FF are concatenated in that order, with each history arranged from newest to oldest, giving the 80-dimensional feature vector φk in Eq. (2).

$$q_k = \log_{10}(1 - x_{D,k} + \varepsilon), \varepsilon = 10^{-6} \quad (1)$$

$$\varphi_k = [q_k \ldots q_{k-19}, L_k \ldots L_{k-19}, Q_{R,k} \ldots Q_{R,k-19}, F_{F,k} \ldots F_{F,k-19}]^T \quad (2)$$

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

The regression target is the increment Δqk = qk+1 − qk. Twenty historical training trajectories are used, including local reflux and reboiler-duty steps, multilevel random inputs, and earlier feed-flow and nominal-operation records. After constructing the history windows, the earliest point in each window must be at least 40 s after the trajectory starts. This yields 8,298 valid regression samples. Each sample is weighted inversely to the number of valid samples in its trajectory, and the weights are rescaled to have a mean of one over the full dataset. Features are normalized using their weighted means, μ, and componentwise standard deviations, sφ, with a standard-deviation floor of 10^{-12}.

$$z_i = \frac{\varphi_i - \mu}{s_\varphi}, \quad \min_{b,\theta} \sum_{i=1}^{N} w_i(\Delta q_i - b - \theta^T z_i)^2 + \alpha_R\|\theta\|^2 \tag{3}$$

Division in Eq. (3) is componentwise. The intercept is the weighted mean of Δq, and the coefficients are obtained from the weighted normal equations with αR = 1. The history length and ridge coefficient were selected using a historical development trajectory from {5, 10, 20} and {0.01, 1, 100}, respectively. The model is frozen throughout the subsequent BO stage and is not reidentified using the present confirmation records. Equations (4) and (5) give the one-step predictor and inverse output transformation, respectively.

$$\hat{q}_{k+1} = q_k + b + \theta^T \frac{\varphi_k - \mu}{s_\varphi} \tag{4}$$

$$\hat{x}_{D,k+j} = 1 - (10^{\hat{q}_{k+j}} - \varepsilon) \tag{5}$$

Multistep predictions are generated by recursively updating the q history. The predictor is an incremental linear regression in q coordinates and has a nonlinear output mapping in mole-fraction coordinates through the exponential inverse transformation. Accordingly, nonlinear predictive control here refers to this specific structure, rather than optimization with a fully nonlinear first-principles process model. Candidates are rejected if a prediction is nonfinite or q lies outside [−12, 0]; predicted xD values are not clipped. This interval is a numerical safeguard and does not provide a complete physical-feasibility or product-quality constraint.

3.2 Historical Multistep Rollout Assessment

Two historical trajectories, A and B, excluded from the training set are used for multistep rollouts starting at 120, 180, 300, and 480 s, with prediction horizons of 20, 60, and 120 s. Each horizon therefore contains eight rollout windows. Table 2 reports the arithmetic mean of the window-specific root mean square errors (RMSEs), alongside an independently identified model without FF and a persistence predictor. The rollouts use the subsequent inputs recorded in the historical trajectories. They assess prediction conditional on those inputs, rather than online disturbance-preview error. Both trajectories belong to a previously studied feed-flow disturbance family and therefore do not establish generalization to an independent disturbance category.

Table 2. Mean window RMSE in historical rollouts

Horizon (s)	Model with FF	Model without FF	Persistence
20	1.103e-07	4.874e-07	1.595e-06
60	5.270e-07	1.721e-06	4.347e-06
120	2.023e-06	3.936e-06	8.968e-06

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

At a 60 s horizon, the model including FF has a mean window RMSE of 5.270×10^{-7}, compared with 1.721×10^{-6} for the model without FF and 4.347×10^{-6} for persistence prediction. This local assessment supports including FF and using a 60 s online prediction horizon. Errors increase with the horizon and are larger for stronger historical disturbances. The model is consequently used only within the local operating region investigated.

3.3 Three-block Receding-Horizon Optimization

The prediction horizon comprises thirty 2 s samples. Future reflux is parameterized by three constant 20 s control blocks. The normalized increment of block b is $\delta b = (Lb - Lb{-}1)/sL$, with candidates drawn from $\{-1, -0.9, \ldots, 0, \ldots, 0.9, 1\}$. There are $21^3 = 9{,}261$ combinations before constraint screening. All candidates use the same model and future-disturbance assumptions. Equation (6) defines the online cost, subject to the operating constraints in Eq. (7).

$$J_{in} = \frac{1}{30}\sum_{j=1}^{30}\left(\frac{\hat{x}_{D,k+j} - r}{10^{-5}}\right)^2 + \lambda \sum_{b=1}^{3} \delta_b{}^2 \tag{6}$$

$$L_{min} \le L_b \le L_{max}, \; |L_b - L_{b-1}| \le s_L \tag{7}$$

The parameter λ penalizes control moves, and $r = 0.999005$. As the q recursion is affine in the future values of the flux blocks, the hold prediction and individual block responses can be computed first, followed by batch evaluation of the discrete combinations. Only the first move of the selected plan is executed, and optimization is repeated after 20 s. Reflux is held when no candidate improves on the cost of the hold plan. A fixed enumeration order comes in useful to resolve ties and make replay reproducible. Online optimization is explicit about L and its increments. Purity is dealt with by reference headroom and measured from the closed-loop records; $xD \ge 0.999$ is not imposed as a hard prediction constraint.

For feed-flow tests, FSP is used to preview future feed flow. A local discrete response fit gives $a = 0$. Equation (8) therefore uses the current FF at $j = 1$ and FSP for $j \ge 2$, with a^0 evaluated as one. This simplified preview does not imply zero delay in the physical feed-flow response. For the composition-disturbance test, FF is held at its current value in the prediction, and feed mole fractions are not explicit model inputs.

$$\hat{F}_{F,k+j} = F_{SP} + (F_{F,k} - F_{SP})a^{j-1}, j \ge 1 \tag{8}$$

3.4 PI Benchmark

The PI gains are fixed at their previously developed values, $Kp = 4.423143961765136$ and $Ti = 120$ s, with a control interval of $Tc = 20$ s. The PI benchmark uses the same reference, operating limits, rate limits, and common initial condition as predictive control. The normalized error is $ek = (r - xD,k)/10^{-5}$, and the unconstrained reflux command vk is calculated using Eq. (9). The integral state is initialized for bumpless transfer from the current L as $Iinit = (L - Lnom)/sL - Kpek$.

$$v_k = L_{nom} + s_L\left(K_p e_k + I_k + K_p \frac{T_c}{T_i} e_k\right) \tag{9}$$

The requested reflux, Lreq, is the projection of vk onto $[\max(Lmin, Lk - sL), \min(Lmax, Lk + sL)]$. The integral state is updated using Eq. (10), followed by an additional correction of $(Lread - Lreq)/sL$ after command confirmation to provide tracking anti-windup. The PI controller uses neither FSP nor composition feedforward. The PI and predictive controllers therefore have different information in the flow tests, whereas the two predictive-control weight configurations use identical information.

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

$$I_{k+1} = I_k + K_p \frac{T_c}{T_i} e_k + \frac{L_{req} - v_k}{s_L} \quad (10)$$

4. Bayesian Weight Tuning with a Finite Evaluation Budget

4.1 Closed-loop Score and Margin Screening

The outer loop tunes only $\lambda \in [0.01, 0.20]$. A function evaluation means a complete dynamic co-simulation run, not a GP prediction or a physical plant experiment. Each candidate is evaluated in the negative-flow case after restoring the common complete snapshot; weights are not switched within a continuous trajectory. To prevent candidates from benefiting merely by changing their own scoring scale, all runs use the fixed external score in Eq. (11), with a tracking normalization of 10^{-5} and a constant move coefficient of 0.05. The negative-flow evaluation grid is 122:2:1200 s in acquisition wall time. The move term covers the 54 control opportunities at 120:20:1180 s, including zero requested increments when reflux is unchanged.

$$J_{eval} = \frac{1}{N_g}\sum_{i=1}^{N_g}\left(\frac{x_{D,i} - r}{10^{-5}}\right)^2 + \frac{0.05}{54}\sum_{k=1}^{54}\left(\frac{L_{req,k} - L_{before,k}}{s_L}\right)^2 \quad (11)$$

In Eq. (11), $x_{D,i}$ is obtained by linear interpolation of valid samples, and $L_{before,k}$ is the readback immediately before the request. Jin evaluates a future three-block plan, whereas Jeval evaluates a complete executed run. The two costs serve two different purposes. The move term measures requested changes rather than actuator travel of actual moves. Each run scores only when checks of completeness, initial condition, disturbance, operating envelope and write status are passed. Runs incomplete or with uncertain status are not included; NaN and Inf are not entered as GP observations.

In addition to the product threshold of 0.999, BO candidates are subject to a post-run engineering-margin test, $x_{min} \geq 0.999003$, where x_{min} is the minimum xD among the original valid samples. Passing runs receive $J_{BO} = \ln(1 + J_{eval})$. Complete runs that fail the additional margin receive $1000 + [\max(0, 0.999003 - x_{min})/10^{-5}]^2 + \min(J_{eval}, 100)$. This finite-penalty rule was fixed before the first GP acquisition. All observations entered into the GP in this study pass the additional margin. They do not establish or certify a product-quality feasibility boundary.

4.2 GP Surrogate and Expected Improvement

The transformation $u = [\log10(\lambda) - \log10(0.01)]/[\log10(0.20) - \log10(0.01)]$ maps the search domain to [0, 1]. All individual JBO observations are standardized using their mean and sample standard deviation; the two original observations at $\lambda = 0.05$ are retained separately. The GP uses a Matérn 5/2 kernel with length scale $\ell = 0.35$ and unit signal variance, as given in Eq. (12)[7].

$$\kappa(u, u') = \left(1 + \frac{\sqrt{5}\,d}{\ell} + \frac{5\,d^2}{3\,\ell^2}\right)\exp\left(\frac{-\sqrt{5}\,d}{\ell}\right), d = |u - u'| \quad (12)$$

Let y denote the standardized observations, K the kernel matrix, and η^2 the diagonal observation-noise variance. The vector k_u contains the kernel covariances between the query and observed locations. The latent GP posterior mean and variance are given by Eqs. (13) and (14), respectively[7]. The raw JBO noise standard deviation is estimated from the two replicates at 0.05 as approximately 2.1539×10^{-4}, with a floor of 10^{-6}. It is converted to the standardized output scale, and a numerical jitter of 10^{-10} is added. Kernel parameters are held fixed to avoid unstable hyperparameter fitting with very few observations.

$$\mu(u) = k_u^T(K + \eta^2 I)^{-1} y \quad (13)$$

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

$$\sigma^2(u) = \kappa(u,u) - k_u{}^T(K + \eta^2 I)^{-1}k_u \quad (14)$$

The minimization-oriented plug-in EI in Eq. (15) is used. The incumbent, finc, is the minimum latent posterior mean over observed weight locations, and $\xi = 0.01$ is defined in standardized output units. The functions Φ and φ are the standard normal cumulative distribution and density functions, respectively. EI is set to zero when the posterior standard deviation is zero. This criterion uses the latent function posterior and is not exact noisy EI [8]. EI is evaluated at 2,001 logarithmically spaced candidate weights. Candidates within a normalized distance of 0.001 of an observed location are excluded, and the remaining maximum is selected.

$$EI(u) = (f_{inc} - \mu(u) - \xi)\Phi(z) + \sigma(u)\varphi(z), z = \frac{f_{inc} - \mu(u) - \xi}{\sigma(u)} \quad (15)$$

The two existing baseline runs at 0.05 are first reused, and two initial-design runs at 0.02 and 0.10 are added. The first GP therefore contains four observations at three locations. EI selects 0.01, which is evaluated in a new co-simulation run before the GP is updated. The updated GP suggests approximately 0.010045037 as the next candidate. Replay at the 54 fixed historical states of the measured 0.01 trajectory produces identical first moves and complete three-block plans for these two weights. The remaining budget is therefore allocated to new confirmation runs with 0.01 frozen. This decision reflects the limited budget and discrete control resolution; it is not an EI-zero or global-convergence stopping rule. Only one of the maximum three EI-selected evaluations is executed. In Fig. 1, the shaded band is the latent GP posterior mean ±1.96 standard deviations, not the probability of meeting the purity threshold.

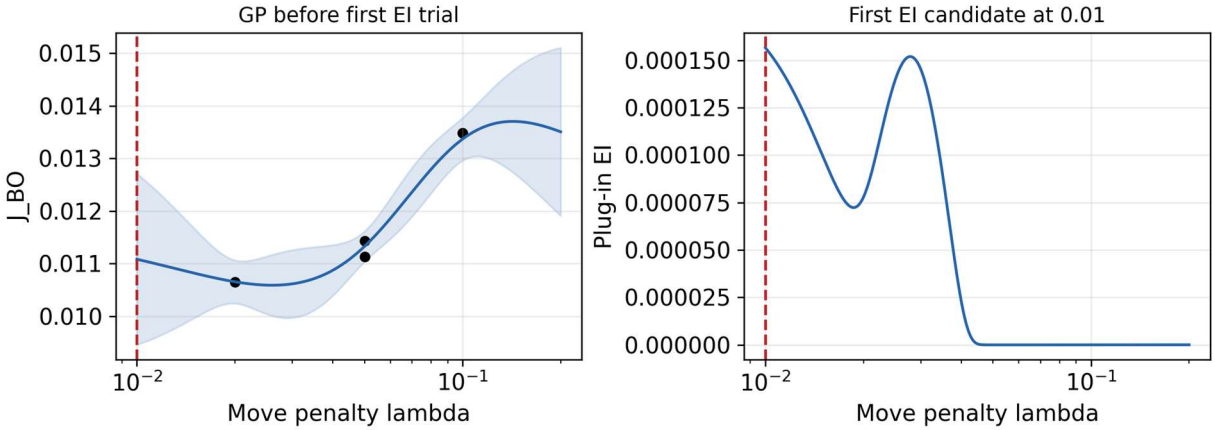


Fig. 1 GP posterior and expected improvement before the first acquisition

The procedure is to freeze predictor and operating envelope, assemble a valid initial observation under the fixed score, compute the GP and EI, choose the candidate run and apply the protocol checks, update the model and then freeze λ for new negative-flow, positive-flow, and composition confirmation runs. No probabilistic constraint surrogate is constructed. The procedure is therefore BO tuning with engineering screening rather than constrained BO with a probabilistic safety guarantee[11,12].

5. Test Design and Evaluation

5.1 Disturbances and Roles of the Data

Table 3 defines the three disturbance cases. The nominal feed mole fractions are acetone 0.44, IPA 0.066, water 0.4937, and hydrogen 0.0003. The composition disturbance reduces acetone by 0.1%

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

relative to its nominal value, to 0.43956, and compensates by increasing water to 0.49414. IPA and hydrogen remain unchanged, so the total mole fraction remains one. In the flow tests, FSP is written automatically once at an acquisition time of 120 s. In the composition test, the two fractions are written sequentially while Aspen is paused, all four fractions are checked, Aspen is resumed, and acquisition is then started.

Table 3. Disturbance cases and control schedules

Case	Disturbance	Control times (s)	Opportunities
Flow decrease	FSP = 76.1507867862 kmol/h	120–1180, every 20 s	54
Flow increase	FSP = 76.3032408138 kmol/h	120–1180, every 20 s	54
Composition decrease	Acetone 0.43956; water 0.49414	40–1180, every 20 s	58

Each run is acquired for 1,200 s, producing 601 samples at a nominal 2 s interval. The 15 valid closed-loop runs comprise six runs with the original $\lambda = 0.05$, two per disturbance case; three PI runs, one per case; three additional BO selection runs, one each at 0.02, 0.10, and the EI-selected 0.01; and three runs after freezing 0.01, comprising a negative-flow repeat and positive-flow and composition confirmations. Startup attempts that fail initial-condition checks and duplicate uploads are not counted as valid runs. Historical identification and rollout records are separate from these 15 closed-loop runs.

The three frozen-parameter confirmation runs are not fed back into BO for additional tuning. Because the disturbance categories and common initial condition had already been exercised during development, these are new runs following parameter freezing, not blind tests on new disturbance categories. The valid runs have 9015 samples in total; within each time series the samples are correlated. There are 15 run-level records; the sample count can't be substituted for statistical replicates.

5.2 Metrics and Interpretation of Time

Quality is evaluated using the minimum xD in the original valid samples and the number of samples below 0.999. Tracking is evaluated by RMSE relative to 0.999005. Move metrics are the mean squared requested increment and the number of reflux writes. The flow cases use the common wall-time grid 122:2:1200 s, whereas the composition case uses 2:2:1200 s. The descriptive Jeval for the composition case has the same form, with the move-term averaging denominator changed to 58. Scores from different cases are not directly compared as measures of control difficulty.

Online control and disturbance scheduling use acquisition wall time. Pointwise equivalence with the internal Aspen simulation clock is not required. In particular, there is an unknown startup interval between resuming Aspen and starting acquisition in the composition tests. Exact simulation settling times aren't reported. The reflux confirmation tolerance is 5×10^{-4} kmol/h; so even a very small request can be confirmed with an earlier readback still within tolerance. The write counts describe controller deployment behaviour not actual valve travel or energy savings.

6. Results and Discussion

6.1 BO selection and Confirmation with the Weight Frozen

Table 4 lists the BO observations and frozen-weight repeat for the negative-flow case. There are other additional initial designs, for example 0.02. Here 0.02 has a lower score than 0.10. The first EI acquisition finds a run for the lowest search bound, 0.01, demonstrating that this point is produced by the surrogate and acquisition criterion rather than prescribed as an initial design. The selection run at

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

0.01 gives Jeval = 0.00831113, and the frozen-weight repeat gives 0.00880024. Both are lower than the two original-weight results.

Table 4. Weight selection and confirmation for the negative-flow case

λ	Data role	Jeval	Minimum xD	Writes
0.05	Original-weight baseline	0.01118722	0.999004185	24
0.05	Original-weight baseline	0.01149528	0.999004185	22
0.02	Initial design	0.01070338	0.999004900	28
0.10	Initial design	0.01357953	0.999003708	16
0.01	EI acquisition	0.00831113	0.999004900	26
0.01	Frozen-weight repeat	0.00880024	0.999004900	30

The mean score of the two runs at 0.05 is 0.01134125. The frozen-weight repeat reduces this mean by 22.41%, whereas the reduction in the selection run is 26.72%. The 22.41% value is reported as the result of the new confirmation run. The frozen-weight repeat has a minimum xD of 0.999004900 and passes the additional 0.999003 margin test. The best observed weight lies at the lower search bound, so the results do not establish 0.01 as a global optimum beyond the investigated domain.

Hereafter, the original-weight predictive controller is denoted PC0, and the predictive controller with frozen λ = 0.01 is denoted BO-PC. Their predictor, horizons, grid, bounds, and disturbance preview are identical; only the move penalty differs.

6.2 Quality Maintenance under the Negative-flow Disturbance

Table 5. Closed-loop results for the negative-flow case

Controller	Minimum xD	End xD	RMSE	Jeval	Writes
PI	0.999001205	0.999005020	1.779e-06	0.033597	54
PC0 first	0.999004185	0.999005973	9.734e-07	0.011187	24
PC0 repeat	0.999004185	0.999005973	9.900e-07	0.011495	22
BO-PC confirmation	0.999004900	0.999005854	8.179e-07	0.008800	30

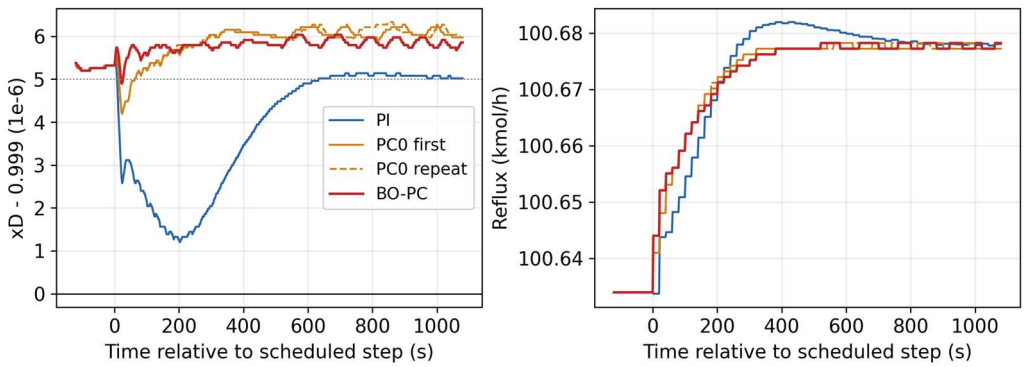


Fig. 2 Top-product purity margin and reflux under the negative-flow disturbance

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

As shown in Fig. 2, the PI quality margin initially decreases substantially after the negative-flow step, reaching a minimum xD of 0.999001205. Both PC0 runs reach 0.999004185, whereas the BO-PC confirmation reaches 0.999004900. All configurations meet the product threshold. BO-PC retains a larger minimum margin in this record and has a tracking RMSE of 8.179×10^{-7}, below the PC0 values of 9.734×10^{-7} and 9.900×10^{-7}. The dotted line represents the reference margin of 5×10^{-6} above the product threshold, and the zero line is the product threshold.

BO-PC writes the reflux 30 times, compared with 24 and 22 writes for PC0 (which imply more active reflux adjustment with a smaller move penalty). Score improvement only reflects the fixed tracking–movement trade-off and does not mean fewer writes or even lower actual operating consumption. Because PI has no FSP feedforward, its comparison with BO-PC reflects the combined effect of the prediction, preview and the tuning configuration. Difference cannot be due to BO only.

6.3 Positive-flow Confirmation and Persistent Excess Purity

Table 6. Closed-loop results for the positive-flow case

Controller	Minimum xD	End xD	RMSE	Jeval	Writes
PI	0.999005198	0.999045670	3.363e-05	11.311265	54
PC0 first	0.999005079	0.999046743	3.438e-05	11.819587	3
PC0 repeat	0.999005079	0.999046743	3.438e-05	11.823109	3
BO-PC confirmation	0.999004960	0.999046743	3.431e-05	11.771740	2

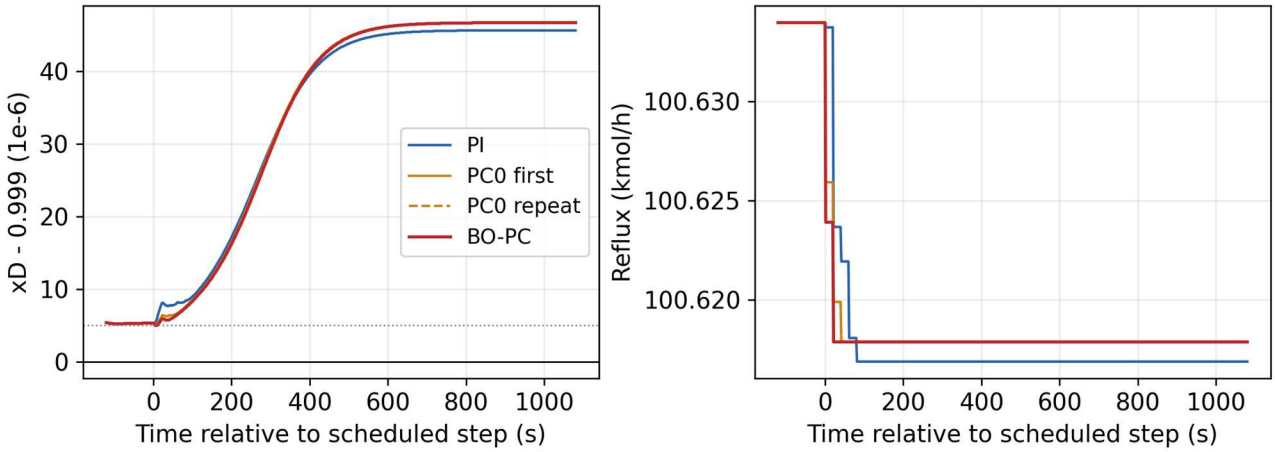


Fig. 3 Top-product purity margin and reflux under the positive-flow disturbance

In the positive-flow case, the BO-PC confirmation has a minimum xD of 0.999004960 and an end value of 0.999046743, identical to the PC0 end value. Its Jeval is 11.771740, approximately 0.42% below the two-run PC0 mean of 11.821348. This change is much smaller than in the negative-flow case. The PI RMSE is 3.363×10^{-5}, slightly smaller than either predictive control configuration. This case therefore also proves sampled quality compliance after freezing of the parameters, but not a claim of clear better performance in the three cases.

Figure 3 shows continued excess purity for all the configurations after reflux reaches its lower limit. The small range for reflux and fixed QR limits further ability to adjust towards the tracking reference. Changing λ alone does not eliminate this deviation within the prescribed envelope of operation. Since $xD \geq 0.999$ is the main objective, excess purity is not considered as a product-quality failure. But the results do not confirm simultaneous precise tracking of the reference and energy-efficient operation.

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

6.4 Confirmation under the Composition Decrease

Table 7. Closed-loop results for the composition-decrease case

Controller	Minimum xD	End xD	RMSE	Jeval	Writes
PI	0.999002576	0.999004960	1.090e-06	0.012490	58
PC0 first	0.999004483	0.999005258	2.773e-07	0.001312	21
PC0 repeat	0.999004364	0.999005318	3.163e-07	0.001526	19
BO-PC confirmation	0.999004483	0.999005198	1.838e-07	0.000984	25

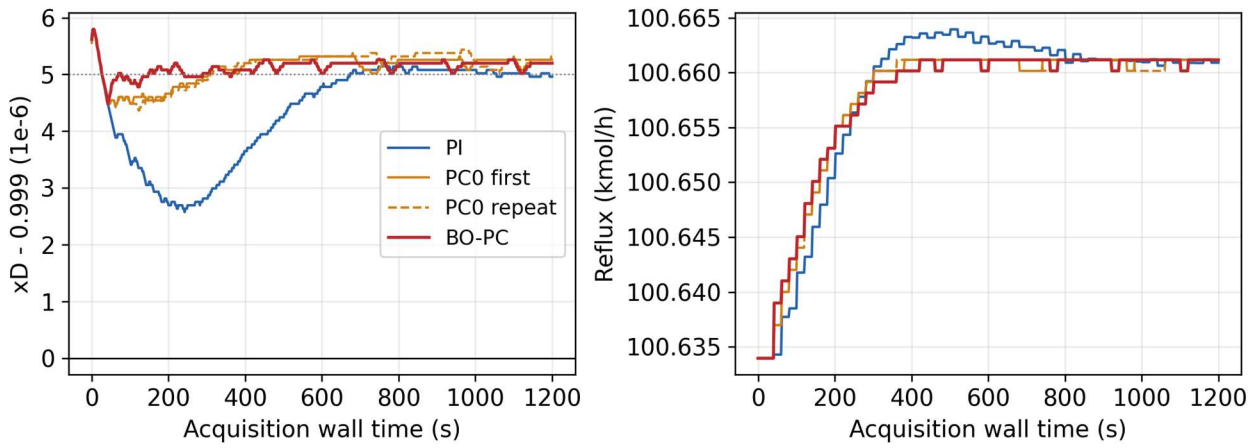


Fig. 4 Top-product purity margin and reflux under the feed-composition decrease

Under the composition decrease, the BO-PC confirmation has a minimum xD of 0.999004483, identical to that in the first PC0 run. The PC0 repeat reaches 0.999004364 and PI reaches 0.999002576. BO-PC ends at 0.999005198, with a tracking RMSE of 1.838×10^{-7}. Its descriptive score of 0.000984287 is 30.64% below the two-run PC0 mean of 0.001419151. All records meet the product threshold. BO-PC writes reflux 25 times, compared with 21 and 19 writes for PC0.

These results demonstrate sampled quality compliance and a lower descriptive tracking error for this local composition disturbance without explicit composition feedforward. However, there is only one frozen-parameter confirmation, and the startup interval for the composition disturbance was not measured precisely. The 30.64% reduction is therefore a comparison of the available records, not a confirmed and repeatable gain in statistics. Minimum quality margin is also unchanged from the first PC0 run.

6.5 Scope of the Evidence and Methodological Limitations

The three frozen-parameter confirmation runs contain 1,803 valid samples, none below 0.999. All 9,015 samples from the 15 valid runs also meet the product threshold. This evidence concerns recorded discrete samples. It does not establish continuous-time quality guarantees between samples, before acquisition starts, or for arbitrary disturbance amplitudes. Product compliance and reference tracking must be evaluated separately: the positive-flow case meets the quality threshold but has a relatively large tracking offset, whereas the negative-flow score improvement is accompanied by more reflux writes.

The number of runs is small and unbalanced: PC0 has two runs per case, whereas PI and frozen-parameter confirmation have one each. Reliable performance confidence intervals and significance conclusions cannot be obtained from this design. The GP noise estimate also relies on only two

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

baseline runs. Executing one EI-selected candidate establishes that a sequential acquisition was implemented, but does not show greater sample efficiency than random search, grid search, or manual tuning. Such a comparison would require optimization baselines with equal evaluation budgets.

The predictor is local and data-based. Although QR is still in the historical feature vector, it is fixed online; the results have not demonstrated co-ordinated control with two manipulated variables. The common snapshot has not been shown to represent the steady state for a full column, and the internal simulation clock is not explicitly synchronized with wall time. Complete Aspen modelling settings, including property method, pressure and liquid holdup, should accompany the dynamic project file so that it can be reproduced on another platform. These limitations do not alter the numerical comparisons of the present records, but restrict their extension to larger disturbances, other initial conditions, and the full process.

7. Conclusion

A ridge incremental predictor with a nonlinear output transformation, a three-block receding-horizon controller under reflux magnitude and rate limits and a GP/EI weight-tuning procedure with post-run purity-margin screening were developed for overhead quality maintenance for the acetone purification column T2. With the predictor and operating envelope fixed, BO selected $\lambda = 0.01$, which was investigated in new runs after freezing the parameters.

In the negative-flow repeat, BO-PC had a minimum value of the overhead acetone mole fraction 0.999004900 and decreased the fixed performance score by 22.41% over the mean of two runs with the original $\lambda = 0.05$. The positive-flow and composition confirmations achieved minimum values of 0.999004960 and 0.999004483, respectively, with all valid samples being >0.999. The improvement in the positive-flow score is only about 0.42%, and the excess purity still exists. 30.64% score decrease of the composition case is a descriptive result from one case confirmation. Thus, the evidence confirms local sampled quality compliance and improved aggregate performance in the negative-flow case but does not indicate global optimality, universality or guaranteed continuous time hard quality.

More work will improve simulation time synchronization and provide the full project settings needed for reproduction; more replicate runs will be necessary, and tuning method comparisons should be made under an equal evaluation budget. Further experiments should test other internal initial conditions and larger disturbances, increasing the manipulated-variable set as appropriate for joint quality and energy goals.

References

[1] Skogestad, S. (1997). Dynamics and control of distillation columns-A critical survey. *Modeling, Identification and Control*, 18(3), 177–217. https://doi.org/10.4173/mic.1997.3.1

[2] Rawlings, J. B., Mayne, D. Q., & Diehl, M. M. (2017). *Model predictive control: Theory, computation, and design* (2nd ed.). Nob Hill Publishing. https://sites.engineering.ucsb.edu/~jbraw/mpc/

[3] Hewing, L., Wabersich, K. P., Menner, M., et al. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3, 269–296. https://doi.org/10.1146/annurev-control-090419-075625

[4] Chen, S., & Billings, S. A. (1989). Representations of non-linear systems: The NARMAX model. *International Journal of Control*, 49(3), 1013–1032. https://doi.org/10.1080/00207178908559683

[5] Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67. https://doi.org/10.1080/00401706.1970.10488634

[6] Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13, 455–492. https://doi.org/10.1023/A:1008306431147

[7] Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press. https://doi.org/10.7551/mitpress/3206.001.0001

ISSN: 2414-1895 DOI: 10.6919/ICJE.202611_12(11).0005

[8] Frazier, P. I. (2018). A tutorial on Bayesian optimization. arXiv preprint arXiv:1807.02811. https://doi.org/10.48550/arXiv.1807.02811

[9] Marco, A., Hennig, P., Bohg, J., et al. (2016). Automatic LQR tuning based on Gaussian process global optimization. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation* (pp. 270–277). IEEE. https://doi.org/10.1109/ICRA.2016.7487144

[10] Kudva, A., Huynh, M. T., Mesbah, A., et al. (2024). Efficient performance-based MPC tuning in high dimensions using Bayesian optimization over sparse subspaces. *IFAC-PapersOnLine*, 58(14), 458–463. https://doi.org/10.1016/j.ifacol.2024.08.379

[11] Xu, W., Jones, C. N., Svetozarevic, B., et al. (2024). Violation-aware contextual Bayesian optimization for controller performance optimization with unmodeled constraints. *Journal of Process Control*, 138, 103212. https://doi.org/10.1016/j.jprocont.2024.103212

[12] Gardner, J. R., Kusner, M. J., Xu, Z. E., Weinberger, K. Q., & Cunningham, J. P. (2014). Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning* (Vol. 32, No. 2, pp. 937–945). PMLR. https://proceedings.mlr.press/v32/gardner14.html