Reliable Backtesting for Continuous Industrial Optimization via Sequential Simulation: A Gas-Fired Ceramic Kiln Case Study

Silvia Ventura-Cabrejas[1][0009-0009-9618-613X] silvia.ventura@urjc.es, Dayrelis Mena-Torres[2][0000-0001-8610-5791] dmentor@upo.es, Luis García Marín[3] lgarcia1@us.es, and Francisco García Flores[4] francisco.garcia@oga.ai

[1] Rey Juan Carlos University, Madrid, Spain
[2] Pablo de Olavide University, Seville, Spain
[3] Seville University, Seville, Spain
[4] Product Manager, OGA, Seville, Spain

Abstract. Offline evaluation of optimization policies for continuous industrial processes becomes unreliable when each historical timestamp is optimized independently, because a counterfactual control action should modify the plant state observed at later times. This paper studies this issue in a gas-fired mono-layer ceramic kiln, where turbine frequency is adjusted to preserve a target pressure profile while reducing gas consumption. We propose a finite-horizon sequential simulator that propagates a compact process state through auxiliary transition models, reconstructs future optimizer inputs, and enables temporally coherent closed-loop backtesting. The framework combines objective models for gas and pressure compliance with a two-stage constrained optimizer that respects industrial actuation limits. Using the proposed sequential simulation, the optimized policy reached baseline-referenced improvements, computed against a predefined internal improvement threshold, of 20.61% in total gas consumption, 15.17% in pressure-profile compliance, and 38.70% in optimal minutes on the continuity-constrained subset.

Keywords: industrial optimization; backtesting; sequential simulation; ceramic kiln; predictive control; continuous processes.

1 Introduction

The offline evaluation of optimization policies is a central methodological issue in sequential decision problems, because the quantity to be estimated is counterfactual: the available trajectories were generated by a behavior policy that differs from the policy to be evaluated [1,2]. This difficulty is especially relevant in

industrial settings, where online experimentation may be costly, operationally disruptive, or undesirable from a safety perspective [1,2]. In parallel, recent reviews on process control have documented growing interest in reinforcement learning and related sequential optimization methods for process industries, while also emphasizing practical limitations associated with safety, sample efficiency, and model mismatch [3,4]. These considerations are directly relevant in continuous thermal processes, where control decisions affect future operating conditions and therefore also affect the future inputs seen by the optimization system.

Prior work nevertheless leaves a specific gap at the evaluation layer. The literature on offline policy evaluation and offline reinforcement learning establishes the counterfactual nature of the problem and shows that the reliability of historical-policy assessment depends on factors such as distribution shift, horizon length, and modeling assumptions [1,2]. In industrial process control, offline learning has already been used to improve decision policies from historical data without unsafe exploration [5], but this line of work focuses mainly on policy learning rather than on the validity of historical backtesting under repeated counterfactual actions. A related literature on digital twins in process engineering describes data-driven surrogates for monitoring, optimization, and control [6], but does not directly address the narrower question considered here: how to perform a temporally coherent offline backtest when future optimizer inputs should depend on previously selected counterfactual controls.

The kiln application further motivates this gap because ceramic-process optimization is inherently multi-criteria and strongly conditioned by changing production regimes [7,8,9]. Existing work focuses mainly on process modeling and operational optimization, leaving open the question of how to estimate closed-loop policy performance from historical data in a way that remains consistent with the dynamics of a continuous industrial process. This paper addresses that question in a gas-fired mono-layer ceramic kiln by focusing on backtesting methodology rather than on policy learning alone. Starting from a one-step optimization framework for gas consumption and pressure-profile compliance, we propose a finite-horizon sequential simulator that propagates a compact process state and reconstructs future optimizer inputs.

2 Industrial Context and Problem Formulation

2.1 Continuous process and observed state

The case study examines a gas-fired mono-layer ceramic kiln running in continuous production. The process is tracked via several variable groups-temperatures at multiple kiln points, pressures, gas flow and consumption, production indicators, and turbine and motor frequencies-sampled at short, fixed intervals to describe the current operating state. Because the kiln configuration changes with the product being fired, the system lacks a stationary regime that can be predicted reliably over long horizons from historical data alone. Real-time state information

is therefore crucial, and the optimization policy must adapt to the kiln's current condition rather than depend on long-term open-loop forecasts.

2.2 Operational objectives

Let $x_t \in \mathbb{R}^d$ denote the feature vector built at sampling instant t from the current multivariate observation and a short recent history. The first objective is to preserve a target pressure profile along the kiln. Consider an ordered set of pressure sensing locations $\mathcal{K}^{\uparrow} = \{k_1, \ldots, k_r\}$ and a downstream location $k_{\downarrow}$ at which the profile must decrease after reaching its peak. Because instantaneous pressure measurements are affected by fluid-dynamic fluctuations, the profile is defined on smoothed pressures. For a smoothing window of length w_p, let

$$\bar{p}_{k,t}^{(w_p)} = \frac{1}{w_p} \sum_{\tau=t-w_p+1}^{t} p_{k,\tau}. \tag{1}$$

The desired pressure pattern can then be written as

$$\bar{p}_{k_1,t}^{(w_p)} < \bar{p}_{k_2,t}^{(w_p)} < \cdots < \bar{p}_{k_r,t}^{(w_p)}, \qquad \bar{p}_{k_r,t}^{(w_p)} > \bar{p}_{k_{\downarrow},t}^{(w_p)}. \tag{2}$$

This is the primary objective of the control problem because it encodes the operating conditions associated with stable firing quality.

The second objective is to reduce gas consumption over a short predictive horizon (H_g):

$$C_t^{(H_g)} = \sum_{j=1}^{H_g} c_{t+j}. \tag{3}$$

The horizon is chosen to remain predictable while still capturing the effect of the manipulated action on gas consumption.

2.3 Action space and real-time decision setting

The manipulated variable is the frequency of one of the turbines, denoted by f_t. At each sampling instant, the controller selects a new value subject to operational and safety constraints. The admissible domain is represented as

$$\mathcal{F} = \{f \in \mathbb{Z} : f_{\min} \leq f \leq f_{\max}\}, \tag{4}$$

and consecutive decisions must also satisfy a rate-of-change restriction

$$|f_t - f_{t-1}| \leq \Delta_{\max}. \tag{5}$$

Here $f_{\min}$, $f_{\max}$, and $\Delta_{\max}$ are determined from historical operating practice, and the decision must be computed within a strict real-time compatible with deployment in production.

2.4 Predictive modeling of the objectives

The optimization relies on two predictive models evaluated under a candidate action $f \in \mathcal{F}$.

The pressure model estimates the probability that the smoothed pressure profile in (2) will be satisfied over the predictive window:

$$\hat{p}_t(f) = g(x_t, f), \tag{6}$$

where g is a supervised learning model. A binary feasibility indicator is then defined by thresholding that probability at an operating threshold τ_p,

$$\hat{y}_t(f) = \mathbb{I}[\hat{p}_t(f) \geq \tau_p]\,. \tag{7}$$

The gas model estimates the cumulative gas consumed over the predictive horizon,

$$\hat{C}_t(f) = h(x_t, f) \approx C_t^{(H_g)}, \tag{8}$$

where h is a regression model trained to approximate the short-horizon cumulative gas objective in (3).

2.5 Two-stage constrained optimization

The original control logic is best described as a two-stage optimization scheme with a fallback rule. Let

$$\mathcal{N}(f_{t-1}) = \{f \in \mathcal{F} : |f - f_{t-1}| \leq \Delta_{\max}\} \tag{9}$$

be the locally admissible neighborhood around the previous action.

Stage 1: globally preferred operating point. Ignoring the local rate limit, the algorithm searches the full action domain for the gas-minimizing setting under a large penalty (K_1) on pressure infeasibility, thereby defining a globally preferred operating point.

$$f_t^{\text{glob}} = \arg\min_{f \in \mathcal{F}} \left[\hat{C}_t(f) + K_1\big(1 - \hat{y}_t(f)\big)\right]. \tag{10}$$

Stage 2: locally feasible implementation. The second stage selects the locally admissible action closest to that point under the rate-of-change limit:

$$f_t^{\text{loc}} = \arg\min_{f \in \mathcal{N}(f_{t-1})} \left[|f - f_t^{\text{glob}}| + K_2\big(1 - \hat{y}_t(f)\big)\right]. \tag{11}$$

The role of K_2 is analogous to that of K_1, but now in the restricted neighborhood.

Fallback toward maximum compliance probability. If no local action is predicted to satisfy the pressure constraint, the controller moves toward the action with maximum compliance probability whenever the predicted gain exceeds the relevance threshold δ.

$$f_t^{\max} = \arg\max_{f \in \mathcal{F}} \hat{p}_t(f). \tag{12}$$

Whenever the improvement in predicted compliance probability exceeds a relevance threshold δ,

$$\big(\hat{p}_t(f_t^{\max}) - \hat{p}_t(f_t^{\mathrm{loc}})\big)\big(1 - \hat{y}_t(f_t^{\mathrm{loc}})\big) > \delta, \tag{13}$$

the final action is replaced by the locally admissible setting closest to $f_t^{\max}$:

$$f_t = \arg\min_{f \in \mathcal{N}(f_{t-1})} |f - f_t^{\max}|. \tag{14}$$

Otherwise, the implemented action is simply $f_t = f_t^{\mathrm{loc}}$. This additional rule reflects the original practical reasoning: when the local neighborhood cannot make the process compliant, it is preferable to move toward the action that offers the best chance of recovering compliance, provided that the gain is materially meaningful.

3 From Naive Backtesting to Sequential Simulation

In continuous industrial optimization, the relevant evaluation target is not a single isolated control decision, but the medium-term behavior induced by applying the optimization logic repeatedly over time. From the perspective of offline policy evaluation, the quantity of interest is therefore inherently counterfactual and sequential: what matters is not only the immediate predicted effect of an action at the current instant, but also how that action changes the future trajectory of the process, the future states observed by the controller, and the subsequent decisions that the optimizer will take under those evolved conditions [1]. In the kiln case study, this distinction is especially important because turbine-frequency adjustments do not only affect instantaneous objective values; they also modify future operating conditions, such as thermal and flow-related variables, which later re-enter the optimization problem through the engineered features used by the predictive models. As a result, a meaningful offline assessment must capture not just immediate action quality, but the sequential interaction between control actions, state evolution, and repeated re-optimization.

This requirement exposes a structural limitation of naive historical replay for continuous-process backtesting. If, at time t, one evaluates a counterfactual optimized action but then continues the rollout using the logged historical measurements at time $t + 1$, the resulting trajectory mixes two incompatible worlds: a counterfactual intervention at the current step and a factual continuation generated under different past decisions. Such a procedure may approximate the local one-step effect of changing the manipulated variable, but it cannot represent

the state-steering effect that arises when optimized actions are applied recursively over a finite horizon. For this reason, preserving actuator feasibility or one-step realism alone is insufficient for credible offline evaluation in this setting; the backtesting framework must also include a mechanism to propagate the process state under the optimized decisions, rather than simply reusing the observed future from the logged data [1,3].

3.1 Approximate closed-loop evaluation

The quantity of practical interest is a finite-horizon closed-loop evaluation in which each decision affects the optimizer inputs seen at subsequent steps. If $\tilde{z}_{t+h}$ denotes the counterfactual state that would arise under repeated optimized actions, the relevant trajectory satisfies

$$\tilde{z}_{t+h+1} = \Phi(\tilde{z}_{t+h}, f_{t+h}, \xi_{t+h}), \qquad f_{t+h} = \mu(\tilde{x}_{t+h}, f_{t+h-1}), \qquad \tilde{x}_{t+h} = \Gamma(\tilde{z}_{t+h}). \tag{15}$$

Because the true transition law Φ is unknown, the counterfactual trajectory cannot be reconstructed exactly from historical data alone. The proposed remedy is therefore not exact replay, but sequential simulation. A learned transition operator $\hat{\Phi}$ is used to propagate a compact simulated state,

$$\hat{z}_{t+h+1} = \hat{\Phi}(\hat{z}_{t+h}, f_{t+h}), \qquad \hat{x}_{t+h} = \Gamma(\hat{z}_{t+h}), \tag{16}$$

so that each optimization step is evaluated on inputs that are consistent with the previously chosen optimized actions.

The corresponding approximate sequential evaluation is then

$$\hat{J}_{\text{seq}} = \frac{1}{H} \sum_{h=0}^{H-1} \ell(\hat{x}_{t+h}, f_{t+h}), \tag{17}$$

where H is a finite rollout horizon selected to keep simulation error within acceptable bounds.

3.2 Methodological implication

The issue with naive historical backtesting is not simply that it uses a short horizon, but that it evaluates optimized actions against a future state trajectory that those actions did not generate. In a production environment, the optimizer would always act on a state already altered by previous control moves. A credible offline backtest must preserve that dependence, at least approximately, by propagating the subset of process variables required to reconstruct the next optimizer input.

For this reason, the methodology proposed in this paper introduces auxiliary transition models whose sole purpose is to maintain temporal coherence across optimization steps. Their role is not to build a perfect digital twin of the kiln, but to provide a sufficiently faithful finite-horizon approximation of the counterfactual closed-loop evolution needed for reliable backtesting.

3.3 Data preparation and operational subsets

Industrial historian streams were synchronized and cleaned by removing communication failures, no-production periods, anomalous intervals, incompatible firing-curve segments, and transition regions. Three nested subsets were then defined: training minutes, reliable minutes satisfying the target pressure profile, and optimal minutes corresponding to reliable records with favorable future gas consumption within product context. A relaxed admissibility criterion was later used to enlarge the set of optimizable minutes.

3.4 One-step objective models

The single-step optimization uses two predictive models: a regression model to estimate the average gas consumption over the next five minutes in normal cubic metres (Nm^3), and a classification model to estimate the probability of pressure-profile compliance over the same five-minute horizon for each operating minute. Both models use engineered features derived from synchronized plant measurements, including temporal encodings, lagged summaries, rolling descriptors, and selected interaction terms, with Boruta feature selection using LightGBM as the base estimator to reduce dimensionality and excessive collinearity.

For the gas model, the train/test partitioning was stratified by product type so that the observed operating contexts were consistently represented in both subsets. The resulting model achieved MAE = 0.2218, MAPE = 0.9463%, RMSE = 0.3572, and $R^2 = 0.9712$. Recent gas-flow information, recent production-rate summaries, the product descriptor, and the manipulated turbine frequency were among the most influential predictors.

For the pressure-profile model, the training subset was balanced between compliant and non-compliant examples. Again, multiple candidate algorithms were explored and the final choice was guided by the F1 score and the recall of the negative class, because missing pressure-profile violations is more critical than maximizing average accuracy. The chosen model achieved ACC = 0.6885, BACC = 0.7184, PREC = 0.7596, REC = 0.6836, F1 = 0.7022, and AUC = 0.8017. The most influential features were recent gas-flow summaries, recent temperature summaries, and the turbine frequency itself, which ranked among the leading predictors for pressure feasibility.

3.5 Sequential simulation by auxiliary models

To repair the mismatch identified in Section 3, offline evaluation was reformulated as a finite-horizon sequential simulator. As summarized in Fig. 1, the framework starts from a compact simulated state z_t, reconstructs the engineered features required by the objective models, applies the two-stage optimizer to select the turbine frequency f_t, and propagates the next state z_{t+1} through a bank of auxiliary one-step transition models.

The auxiliary-model bank excludes the manipulated turbine frequency, which is directly imposed by the optimizer, and externally scheduled production descriptors, which cannot be inferred from process dynamics alone.The remaining temperature, gas-flow, and pressure variables were selected under two principles: they had to be predictable one step ahead using only information available at time t, and the propagated state had to remain as compact as possible. This compactness criterion is important because enlarging the propagated state would require many more auxiliary models and would also amplify recursive error accumulation over the rollout. In practice, the best trade-off was to predict each next-step variable from the current compact state and reconstruct the additional rolling and lagged descriptors through the same feature-engineering pipeline used during objective-model training.

The same trade-off guided model-development decisions. Because the framework requires many auxiliary models, exhaustive hyperparameter tuning for each of them would have imposed a disproportionate computational burden. Standard configurations were therefore preferred whenever they provided adequate predictive quality, which is also consistent with the real-time constraints of the intended deployment setting.

This framework behaves differently in deployment and in offline backtesting. In real production, fresh sensor measurements arrive at every step, so the state can be updated directly from the plant. In offline backtesting, by contrast, there are no counterfactual observations after an optimized intervention, so the future state must be propagated through learned transition models. As the rollout length increases, this recursive construction becomes progressively more vulnerable to drift and also reduces the number of valid continuous sequences available for evaluation. For this reason, the proposed methodology adopts a deliberately finite operational horizon: long enough to capture the short-term effect of sequential decisions, but limited to the range in which iterative simulation remains sufficiently faithful for credible backtesting. In the present case study, this criterion led to a 15-minute rollout horizon.

The auxiliary one-step transition models were evaluated on an 80/20 split over 120,000 training minutes. Overall, next-step predictability was high for most propagated variables: the majority of models achieved R^2 values above 0.95, several variables were predicted with near-perfect fit, and only a small number showed more moderate predictive quality. These results support the use of a compact auxiliary-model bank as a practical mechanism for finite-horizon state propagation in the proposed backtesting framework.

3.6 Validation protocol for the sequential simulator

Before using the sequential simulator of Fig. 1 for policy evaluation, we validated it as a trajectory-generation mechanism independently of the optimization logic. The objective of this validation stage was to determine whether the iterated combination of auxiliary transition models, feature reconstruction, and objective

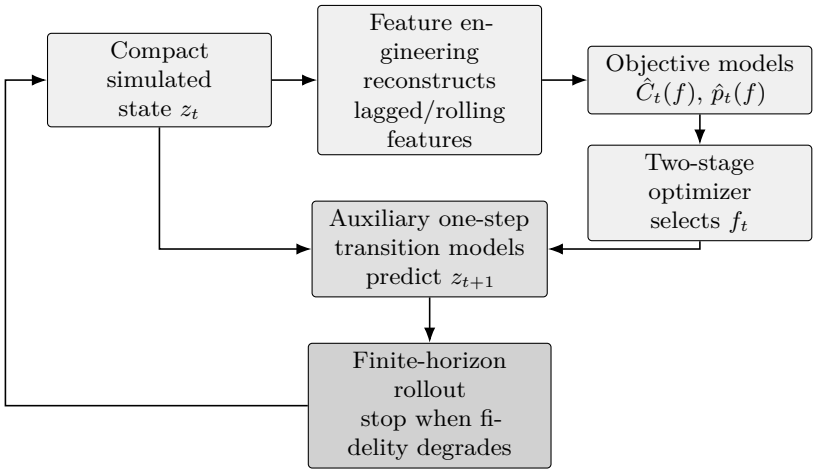


Fig. 1: Finite-horizon sequential simulator for offline backtesting.

models could reproduce the historical short-term evolution of the process with sufficient fidelity. To isolate simulator fidelity from control effects, no optimized intervention was applied during this stage: starting from a real historical minute, the simulator was rolled forward while keeping the observed intervention sequence fixed, and the resulting predicted targets were compared against the recorded historical outcomes.

Validation was performed on optimizable minutes belonging to contiguous sequences long enough to support each candidate rollout. For each horizon, the simulator generated inputs sequentially through the auxiliary-model bank and the original feature-engineering pipeline, and the resulting targets were compared against recorded outcomes using R^2 for cumulative gas prediction and F1 for pressure-profile compliance. The operational horizon was set to the largest rollout satisfying $R^2 > 0.9$ and $\mathrm{F1} > 0.7$ simultaneously. Under this criterion, the maximum reliable horizon was 15 iterations ($R^2 > 0.919$ and $\mathrm{F1} > 0.701$), beyond which drift became too large for credible backtesting.

4 Experimental Evaluation

4.1 Experimental protocol

In total, the dataset comprised over 150,000 minutes. All predictive-model metrics reported for the one-step setting were computed using an 80/20 train/test split applied separately within each temporal block, rather than by selecting training and test samples at random. Several tree-based regression and classification models were evaluated, and XGBoost was selected for the main objective models. One-step optimization results were evaluated over more than 120,000 optimizable minutes

described by the relaxed admissibility criterion. For the sequential simulator, only contiguous sequences long enough to support the selected rollout were retained, leading to around 100,000 comparable minutes. The main regression metric is total or average gas consumption, while the main operational classification metric is the fraction of records satisfying the pressure profile. In addition, the study tracks optimal minutes, defined as reliable records whose future gas consumption lies below the product-specific median. Reported improvements are anonymised, baseline-referenced indices computed against a predefined internal deployment threshold, used as the 0% reference after orienting all metrics so that larger values are better. Thus, positive and negative values indicate margins above or below this threshold, enabling comparison across operational subsets without disclosing absolute consumption values, compliance rates, or targets.

Because the data consist of temporally ordered industrial trajectories, uncertainty should be assessed with block-based resampling over contiguous operational sequences rather than with i.i.d. assumptions.

4.2 Preliminary one-step results and the backtesting problem

Over the approximately 120,000 optimizable minutes, constrained optimization produced a total-gas of -26.01%, a pressure-profile compliance of 3.66%, and an optimal-minutes index of 11.24% (Table 1).

Table 1: One-step optimization results on over 120,000 optimizable minutes, reported as relative margins over an undisclosed internal improvement threshold

Metric	Constrained optimization
Total gas index	-26.01%
Pressure-profile compliance index	3.66%
Optimal minutes index	11.24%

The preliminary analysis also showed that records initially classified as non-reliable were often improved simultaneously in both objectives, indicating that pressure recovery and gas reduction are not necessarily conflicting in this process.

The main limitation of these results is methodological rather than numerical. Each optimized record is evaluated against the original historical continuation, even though the selected control should have modified the observed state at the next instant. Therefore, the one-step backtest provides evidence that the local decision logic is useful, but it does not constitute a faithful closed-loop evaluation.

4.3 Simulation and optimization in continuous processes

The proposed sequential simulator was next used to evaluate the same optimization logic under explicit state propagation. To make the comparison fair, the

previous one-step method and the new sequential method were both evaluated on the same continuity-constrained subset of over 100,000 minutes.

Table 2 shows the key result of the paper. On the common subset, the previous one-step method produced a total-gas index of -25.39%, a pressure-profile compliance index of 3.18%, and an optimal-minutes index of 11.10%. Under the proposed sequential simulation, the corresponding indexed margins increased to 20.61%, 15.17%, and 38.70%, respectively.

Table 2: Comparison on the continuity-constrained subset (over 100,000 minutes), reported as relative margins over an undisclosed internal improvement threshold

Metric	Previous one-step method	Proposed sequential method
Total gas index	-25.39%	20.61%
Pressure-profile compliance index	3.18%	15.17%
Optimal minutes index	11.10%	38.70%

A second important finding is dynamic: cumulative savings increase with the rollout index and then stabilize. This means that the sequential method is not only choosing locally better turbine settings but is also guiding the process toward states that are progressively easier to operate efficiently. When the final iteration is forced to keep the original frequency configuration, part of the accumulated benefit remains. This reveals an *intrinsic saving* term due to state steering rather than to the immediate one-step control action alone.

Finally, when that intrinsic component is removed, the cumulative-saving curve of the sequential method tends toward the savings achieved by the previous one-step method. This behavior is informative: the older backtest essentially captured only the direct local effect of changing the turbine frequency, whereas the new simulator also captures the indirect benefit of driving the process toward more stable and efficient regions of the state space.

5 Conclusions

This paper addressed the offline evaluation of optimization policies in a continuous industrial process and showed that naive one-step historical backtesting is structurally inconsistent with closed-loop operation. When each historical minute is optimized in isolation, the selected action is counterfactual but the subsequent state remains historical, so the evaluation cannot capture how past decisions modify future optimizer inputs.

To repair this mismatch, we proposed a finite-horizon sequential-simulation framework in which a compact subset of process variables is propagated through

auxiliary transition models and used to reconstruct the inputs required by the objective models. The resulting methodology preserves temporal coherence across optimization steps while remaining computationally tractable for industrial use.

The results show that the original one-step optimization already provided useful local improvements in gas consumption and pressure-profile compliance, but that sequential evaluation yields a substantially stronger estimate of the policy value. In the kiln case study, state propagation increased the estimated gas savings and further improved compliance, showing that credible backtesting in continuous processes requires more than isolated historical replay.

Overall, within the validated finite-horizon regime of Section 3.6, state propagation is not a secondary refinement but a necessary component of scientifically defensible offline evaluation in continuous industrial optimization. In this setting, the simulator acts as a short-horizon counterfactual state propagator rather than an unrestricted digital twin, capturing both locally favorable actions and their effect on future operating conditions.

Acknowledgments. This work was developed within the framework of the eIndustry project, in collaboration with the University of Seville, as part of OGA's research and development activities in artificial intelligence and process optimization.

References

1. Lihong Li. A perspective on off-policy evaluation in reinforcement learning. *Frontiers of Computer Science*, 13:911–912, 2019.
2. Cameron Voloshin, Hoang M. Le, Nan Jiang, and Yisong Yue. Empirical study of off-policy policy evaluation for reinforcement learning. In *Advances in Neural Information Processing Systems*, 2021.
3. Rui Nian, Jinfeng Liu, and Biao Huang. A review on reinforcement learning: Introduction and applications in industrial process control. *Computers & Chemical Engineering*, 139:106886, 2020.
4. Joonsoo Park, Hans-Hermann Jung, Jong Woo Kim, and Jong Min Lee. Reinforcement learning for process control: Review and benchmark problems. *International Journal of Control, Automation and Systems*, 23:1–40, 2025.
5. J. Deng et al. Offline reinforcement learning for industrial process control: A case study from steel industry. *Information Sciences*, 632, 2023.
6. L. Peterson et al. Digital twins in process engineering: An overview on computational and numerical methods. *Computers & Chemical Engineering*, 193:108917, 2025.
7. Christos Agrafiotis and Theocharis Tsoutsos. Energy saving technologies in the european ceramic sector: A systematic review. *Applied Thermal Engineering*, 21:1231–1249, 2001.
8. Achilleas L. Arvanitidis, Margaritis Kostoglou, and Michael C. Georgiadis. Modeling, optimization and control of a ceramic tunnel kiln for consistent product quality under changing production demands. *Computers & Chemical Engineering*, 189:108812, 2024.
9. Achilleas L. Arvanitidis, Margaritis Kostoglou, and Michael C. Georgiadis. Optimizing industrial tunnel kiln operations for ceramic roof tile production: A bi-objective approach. *Chemical Engineering Science*, 295:120223, 2024.