Water Science & Technology

Water Science & Technology Vol 00 No 0, 1 doi: 10.2166/wst.2025.139

Model predictive control of anaerobic digestion processes using a long short-term memory network predictor

Andrés Pino Santana [a,b], Santiago Garcia-Gen [a], Laurent Dewasme [b] and Alain Vande Wouwer [b,*]

[a] Departamento de Ingenieria Quimica y Ambiental, Universidad Tecnica Federico Santa Maria, Avenida Espana, 1680, Valparaiso 2390123, Chile
[b] Systems, Estimation, Control and Optimization (SECO), University of Mons, 7000 Mons, Belgium
*Corresponding author. E-mail: alain.vandewouwer@umons.ac.be

APS, 0009-0004-5748-3135; SG-G, 0000-0002-2323-8738; LD, 0000-0002-1688-167X; AVW, 0000-0001-7022-6126

ABSTRACT

This article presents an evaluation of a model predictive control (MPC) strategy that employs long short-term memory (LSTM) networks as internal predictive models for anaerobic digestion (AD) processes. The primary objective was to develop and validate a data-driven control approach utilizing readily available online measurements. The strategy was tested in two simulated AD environments: the simplified Anaerobic Model No. 2 (AM2) and the detailed Anaerobic Digestion Model No. 1 (ADM1). LSTM networks were effectively trained to predict methane flow rates from simulated data, including stochastic disturbances. The integrated LSTM-MPC framework demonstrated robust methane flow rate setpoint tracking and ensured process stability in both environments, even amidst nonlinear operating conditions and influent disturbances. Importantly, the computational requirements remained feasible for real-time applications in these typically slow processes. The findings suggest that the LSTM-MPC strategy is a promising and computationally efficient alternative for controlling AD processes, providing a practical solution compared with traditional mechanistic model-based approaches that often rely on more complex and less accessible measurements.

Key words: anaerobic digestion, long short-term memory, machine learning, model predictive control

HIGHLIGHTS

- LSTM networks provide accurate one-step-ahead predictions of methane flow and adequate multi-step-ahead predictions for MPC.
- The proposed strategy is computationally efficient, making it feasible for real-time control of actual anaerobic digestion processes.
- This approach presents a promising alternative to control methods based on complex mechanistic models.

GRAPHICAL ABSTRACT

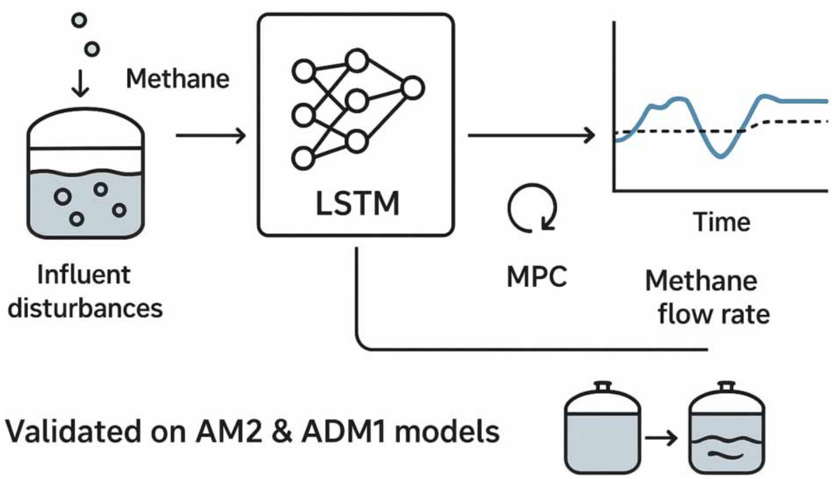

1. INTRODUCTION

Anaerobic digestion (AD) is a widely used biological process for waste treatment and energy production. However, its operation is challenging due to complex, nonlinear, and uncertain dynamics involving multiple microbial populations and reaction pathways (Jimenez *et al.* 2015). This complexity is captured in comprehensive mechanistic models such as the Anaerobic Digestion Model No. 1 (ADM1), but their high number of states and parameters makes them difficult to calibrate and apply for real-time control (Batstone *et al.* 2002; Gaida *et al.* 2012). While simplified mechanistic models like Anaerobic Model No. 2 (AM2) have been developed for control purposes, they are often limited to specific substrates, such as readily biodegradable liquid wastes, and still require significant effort in parameter identification and state estimation, as key variables may not be available for online measurement (Bernard *et al.* 2001; Donoso-Bravo *et al.* 2011; Jamilis *et al.* 2018; Dewasme *et al.* 2019).

The difficulties in mechanistic modeling, coupled with the increasing availability of online process data, have spurred interest in black-box, data-driven approaches (Gupta *et al.* 2023; Rutland *et al.* 2023). Machine learning (ML) techniques, particularly artificial neural networks (ANNs), are effective at predicting biogas flow rate, demonstrating their ability to handle the complex nonlinear dynamics of AD when combined with metaheuristics for variable selection (Beltramo & Hitzmann 2019). Another application of these techniques is the development of software sensors. For instance, ANNs have been used to estimate the evolution of key components from low-cost online measurements, which facilitates process monitoring (Dewasme 2020). Among these, recurrent neural networks (RNNs) such as long short-term memory (LSTM) networks are specifically designed to handle sequential data, making them powerful tools for predicting the dynamic behavior of AD processes (Hochreiter & Schmidhuber 1997; Salamattalab *et al.* 2024). Their utility has been demonstrated for predictive monitoring (McCormick & Villa 2019), and their potential use as internal models for predictive control has been explored for similar chemical processes (Zarzycki & Ławryńczuk 2021).

While accurate prediction is valuable, the ultimate goal is robust process control. Model predictive control (MPC) is an advanced tool that uses a dynamic model to optimize control actions over a future horizon, making it an ideal candidate for managing the slow dynamics and constraints of AD systems (Bolmanis *et al.* 2023). Several studies have demonstrated the effectiveness of nonlinear MPC (NMPC) using simplified mechanistic models, often validated in simulation or at pilot scale (Haugen *et al.* 2014; Kil *et al.* 2017; García-Gen *et al.* 2022). For instance, NMPC has been successfully applied in full-scale plants for demand-driven biogas production (Mauky *et al.* 2016). However, these model-based methods still rely on the availability of a reasonably accurate physics-based model. The literature on combining MPC with purely data-driven, black-box predictors like neural networks (NNs) for AD control is more limited, representing a less explored area in the literature (Gupta *et al.* 2023; Liu *et al.* 2023).

Building upon several of our prior works on modeling, state estimation, and control of the AD process (Giovannini *et al.* 2018; Dewasme *et al.* 2019; Rueda-Escobedo *et al.* 2022; García-Gen *et al.* 2022), this study seeks to bridge this gap. It examines the potential of LSTM networks to serve as effective input–output predictors constructed from online process data. The primary contribution is the evaluation of these data-driven LSTM models when integrated within an MPC scheme, thereby investigating the feasibility and performance of a black-box control strategy for AD systems.

2. METHODS

The proposed control strategy integrates an LSTM network as a data-driven predictor within an MPC structure. The following subsections outline the overall control architecture and provide details on each component.

2.1. Overall control architecture

The control strategy evaluated in this work is based on an MPC scheme where the internal predictive model is a data-driven LSTM network. The overall structure of this closed-loop system is depicted in Figure 1.

The system consists of three main components:

(1) *The plant:* This represents the AD process to be controlled. To evaluate the controller performance under realistic and complex conditions, the plant is emulated in this study using established mechanistic models, which will be detailed in the Results section.

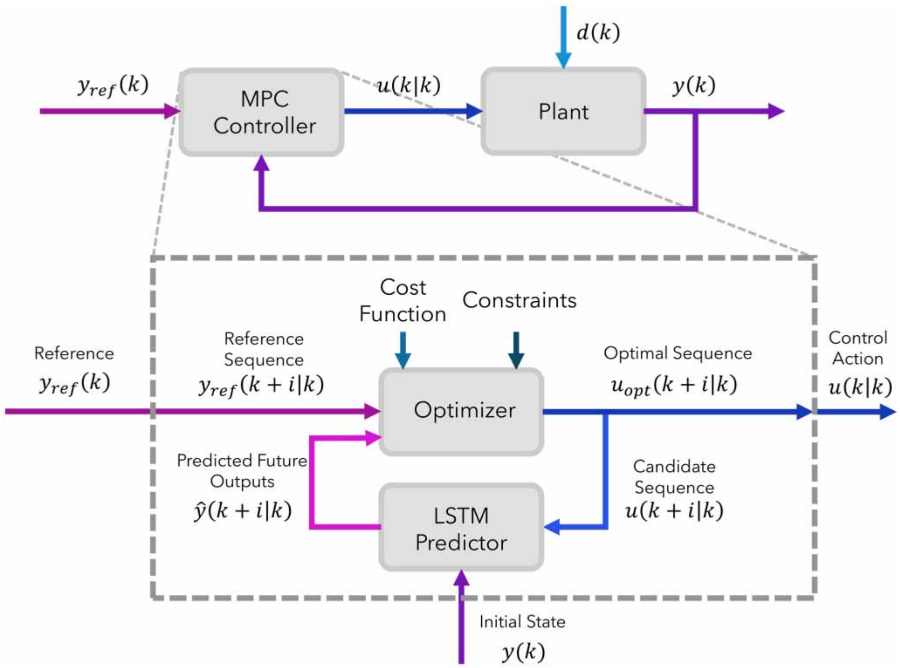

Figure 1 | Schematic of the proposed MPC strategy, integrating a data-driven LSTM network as the internal predictive model.

(2) *The LSTM predictor:* A data-driven model trained on historical input–output data from the plant. At each control interval, it takes a sequence of past manipulated variables (e.g., dilution rate, D) and measured outputs (e.g., methane flow rate, Q_{CH_4}) to predict the future evolution of the outputs.

(3) *The MPC controller:* The core decision-making. It uses the predictions from the LSTM model, the desired reference trajectory (y_{ref}), and operational constraints to solve an online optimization problem. Using a receding horizon, only the initial control action is applied at each time step.

This real-time computational procedure is repeated at each sampling instant, creating a feedback loop where the controller continuously adapts its actions based on the most recent measurements and the model's predictions.

2.2. LSTM predictor model

An LSTM network is a specialized type of RNN with good ability to model complex temporal dependencies in sequential data (Hochreiter & Schmidhuber 1997).

2.2.1. LSTM unit architecture

The core component of an LSTM network is the LSTM unit or cell. Unlike a simple RNN neuron, it contains a memory cell (c_t) and three 'gates' that regulate the flow of information, allowing the network to learn which data in a sequence is important to keep or to forget. The architecture of a single LSTM unit is detailed in Figure 2.

The governing equations for the operations within an LSTM unit at each time step t are:

$$f_t = \sigma_g(W_f x_t + U_f h_{t-1} + b_f) \tag{1}$$

$$i_t = \sigma_g(W_i x_t + U_i h_{t-1} + b_i) \tag{2}$$

$$g_t = \sigma_c(W_g x_t + U_g h_{t-1} + b_g) \tag{3}$$

$$C_t = f_t \odot c_{t-1} + i_t \odot g_t \tag{4}$$

$$o_t = \sigma_g(W_o x_t + U_o h_{t-1} + b_o) \tag{5}$$

$$h_t = o_t \odot \sigma_c(c_t) \tag{6}$$

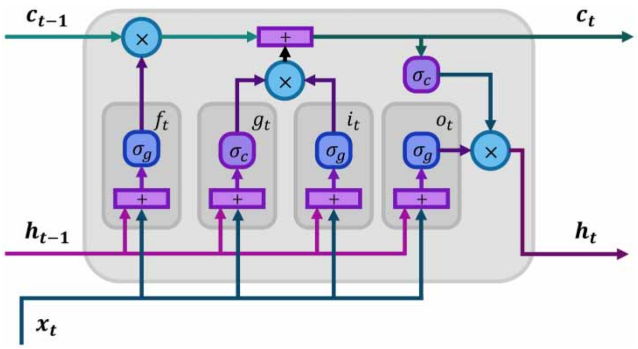

Figure 2 | Detailed architecture of a LSTM unit, illustrating the flow of information through the forget (f_t), candidate cell state (g_t), input (i_t), and output (o_t) gates.

where x_t is the input vector at time t; h_{t-1} is the hidden state from the previous time step; c_{t-1} is the cell state from the previous time step; W and b represent the weight matrices and bias vectors for each gate; σ_g is the sigmoid activation function, σ_c is the tanh activation function; and $\odot$ denotes the element-wise product.

The function of each component is as follows:

- *Forget gate* (f_t): Decides what portion of the long-term memory (c_{t-1}) should be discarded.
- *Input gate* (i_t, g_t): Determines what new information will be stored. The i_t gate selects the information and g_t creates the new candidate values.
- *Cell state* (c_t): Updates the long-term memory by combining the old memory (filtered by f_t) with the new information (filtered by i_t).
- *Output gate* (o_t, h_t): Determines the cell's output. The o_t gate filters the updated cell state (c_t) to produce the hidden state (h_t), which acts as the short-term memory and final output.

2.2.2. Predictor model design and training

The predictor model was designed with an autoregressive structure. It predicts the next output, $\hat{y}(t)$, based on a window of w past values of the output itself (y) and the manipulated variable (u). This relationship is presented by the general equation:

$$\hat{y}(t) = f((u(t-1), \ \ldots, u(t-w), y(t-1), \ \ldots, \ y(t-w)) \quad (7)$$

This function $f(\cdot)$ is approximated by the neural network constructed with a single-layer LSTM network followed by a dense (fully connected) output layer. This architecture, unrolled through time to visualize the flow of sequential information, is shown in Figure 3.

To generate predictions, the model processes data using a sliding window. For a single-step prediction, the model is fed a window with the w most recent historical values. To generate a multi-step prediction over a future horizon, this process is applied iteratively in an autoregressive manner: after predicting the first future output step $\hat{y}(t)$, this predicted value is fed back as an input to predict the next step $\hat{y}(t+1)$.

An important aspect of the model's operation, for both training and prediction, is that it is stateless. This means that the network's internal states (cell state C_t and hidden state h_t) are reset before processing each new input window. Consequently, the memory developed while processing one sequence is not carried over to the next, and each prediction is based solely on the latent representation of the temporal patterns derived from the history contained within that specific window.

The network is trained by minimizing the mean squared error (MSE) loss function between its prediction and the data from the plant.

2.3. MPC framework

MPC is an advanced optimal control method that explicitly uses a dynamic model of the process to compute control actions. At each sampling time k, the MPC algorithm performs the following steps:

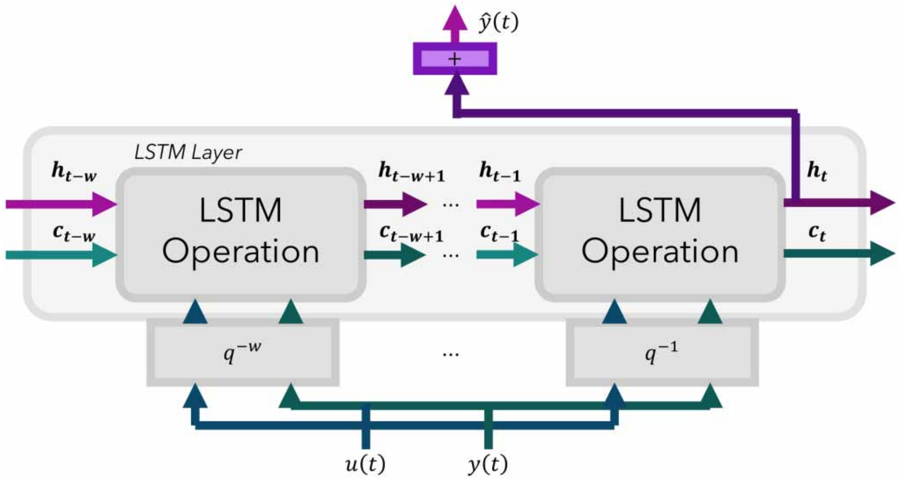

Figure 3 | Unrolled representation of an LSTM layer over a time window. The diagram illustrates the flow of the hidden state (h_t) and cell state (c_t) through time (q^{-w} represents a delay of w sampling times).

(1) *Prediction:* It uses the trained LSTM network to predict the future outputs of the process, $\hat{y}(k+i|k)$, over a *prediction horizon* of N_p steps.
(2) *Optimization:* It calculated an optimal sequence of future control actions $u(k+i|k)$, over a *control horizon* N_c (where $N_c \leq N_p$) by minimizing a cost function, J.
(3) *Implementation:* Only the first control action of the optimal sequence, $u(k|k)$, is applied to the plant. This process is repeated at the next sampling time, $k+1$, using updated measurements from the plant (receding horizon principle).

The cost function mathematically describes the control objectives. A standard quadratic cost function for setpoint tracking, which penalizes both tracking error and control effort, is formulated as

$$J(k) = \sum_{i=1}^{N_p} w_y(\hat{y}(k+i|k) - y_{\text{ref}}(k+i))^2 + \sum_{i=1}^{N_c-1} w_u(\Delta u(k+i|k))^2 \tag{8}$$

where $\hat{y}(k+i|k)$ is the output predicted by the LSTM model at future time $k+i$. $y_{\text{ref}}(k+i)$ is the desired reference value (setpoint). $\Delta u(k+i|k)$ is the change in the control action $(u(k+i|k) - u(k+i-1|k))$. w_y and w_u are non-negative weighting factors that balance the priority between minimizing tracking error and penalizing aggressive control moves.

This optimization is solved subject to operational constraints on the manipulated variables. This makes MPC particularly effective for handling the constrained, nonlinear dynamics of AD processes.

3. RESULTS AND DISCUSSION

This section evaluates the proposed control strategy. The evaluation begins with the simulation setup, followed by a validation of the LSTM predictors, and concludes with an analysis of the closed-loop performance of the integrated LSTM-MPC strategy.

3.1. Simulation setup and process emulation

To assess the performance of the proposed control strategy, the AD process was emulated using two established mechanistic models of varying complexity: the Anaerobic Model 2 (AM2) and the Anaerobic Digestion Model No. 1 (ADM1). These models served as virtual 'plants' to test the controller under diverse and challenging conditions.

The system emulated with the AM2 model was intentionally operated in a highly nonlinear region, close to the operational limits, where reactor washout could occur, thereby creating a significant control challenge. The training data was generated by applying six-step changes to the dilution rate (D) within this critical range. The duration of each step was approximately 4.5

days, and the data were recorded at a sampling rate of approximately 1.75 h. These parameters have been selected to accurately capture the process dynamics.

For the ADM1 model, which represents a more comprehensive and high-fidelity picture of the process, the training data was generated by applying 10-step changes to the inlet flow rate (Q_{in}) across a wide operational range, with each step lasting approximately 14 days. The sampling rate for the ADM1 simulation was 6 h. The adaptation of the temporal parameters is related to the different dynamics of the emulated processes (the two models have not been established for the same effluents, and ADM1 includes a hydrolysis step, which contributes to the slower dynamics).

In both emulation scenarios, to mimic the natural variability of real-world influents, stochastic disturbances were introduced. An AutoRegressive process of order 1 (AR(1)), with parameters $\varphi = 0.95$ and $\sigma = 1\%$ for AM2 or $\sigma = 2.5\%$ for ADM1, was applied to the inlet substrate composition, causing it to fluctuate around a nominal value. This ensures that the LSTM predictors were trained and operated on data that includes both deterministic changes (step inputs) and random noise. The controller's closed-loop performance was then evaluated under these same stochastic conditions.

The prediction horizon (N_p) and the control horizon (N_c) of the MPC controller are both set to 20 steps ahead. The output weight (w_y) is set to 1 for both models. The input weight (w_u), however, is different as it must take into account the distinct nature of the manipulated variable in the two process emulators. Indeed, AM2 considers the manipulation of the dilution rate (D), while ADM1 considers the feed flow rate (Q_{in}). This is reflected in the input weight set to 10,000 for the AM2 emulator and 25 for the ADM1 emulator.

3.2. LSTM predictor performance

Two separate LSTM networks were trained using the data generated from the AM2 and ADM1 emulators, respectively. The network for the AM2 model utilized a window size of 10 past data points (roughly corresponding to the dominant time constant of the plant) and 15 hidden units. For the ADM1 model, window sizes of 10 (again corresponding to the dominant time constant with the adapted sampling interval) and 20 units were used. The performance of these predictors was evaluated through a series of tests on unseen data.

3.2.1. Model training and error analysis

The training progress for both predictors is illustrated in Figure 4. The loss curves show a rapid decrease in MSE for both the training and validation sets, converging to a low, stable value. This indicates that the network learned the underlying process dynamics without significant overfitting.

To evaluate the predictor's long-range accuracy, the growth of the prediction error over an extended prediction horizon was analyzed (Figure 5). The relative root mean squared error (RMSE) increases with the length of the horizon as prediction

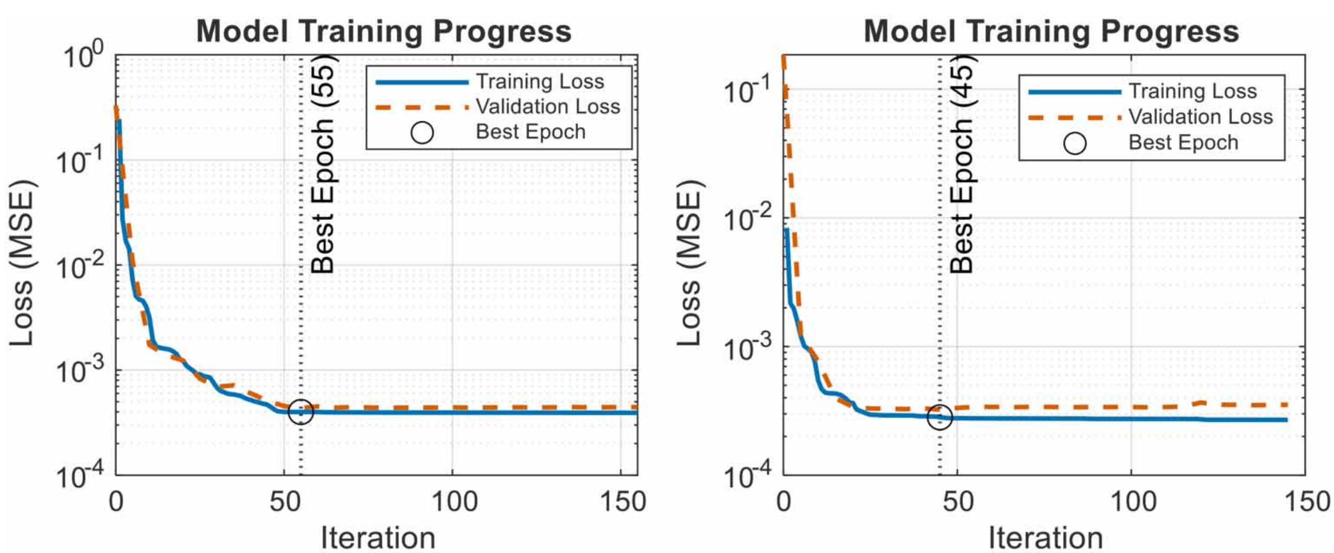

Figure 4 | Training and validation loss curves for AM2 (left) and ADM1 (right) predictors.

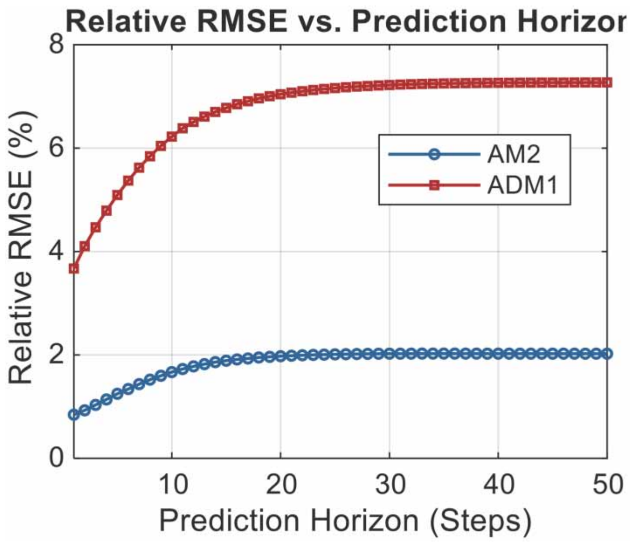

Figure 5 | Multi-step relative prediction error (relative RMSE) as a function of the prediction horizon for AM2 and ADM1 predictors.

errors from earlier steps are propagated to subsequent ones. However, even without real-time measurements, the model's predictions remain stable, a behavior likely attributable to the guiding influence of the sequence of manipulated inputs, which prevents drastic divergence. This analysis provides insights into the predictor's multi-step behavior and highlights the trade-off between a long-term prediction horizon for better planning and a shorter one for higher accuracy. This figure also reveals that the accuracy of the LSTM predictor is better for the case study relative to AM2, but this case study uses a lower level of noise than ADM1, and we observed that the noise level is influential. Moreover, the evolution of the error is similar in both cases.

3.2.2. Evaluation of prediction modes

To assess the suitability of the predictors for MPC, their performance was evaluated in different forecasting modes.

First, the AM2 predictor's behavior was tested in both one-step-ahead and continuous multi-step-ahead modes (Figure 6, top plot). The one-step-ahead predictions (green line), which use process data at each time step, demonstrate high accuracy. The continuous, multi-step predictions (orange line), which are initialized with a window of process data before operating recursively, demonstrate that the predictions stay close to the actual trajectory, capturing the general trend of the process. This behavior is a crucial requirement for use within a control framework.

Second, to most accurately simulate the predictor's function within the MPC's receding horizon, a segmented multi-step prediction test was performed (Figure 6, bottom plot). In this test, the predictor is repeatedly initialized with a window of true historical data to forecast the next N_p steps. This confirms the LSTM's ability to generate reliable short-term forecasts when anchored with real process measurements. While some deviations are noticeable, particularly when the model's strong learned correlation between input–output dynamics causes it to overlook more subtle process behaviors momentarily, the forecasts seem to be sufficiently accurate for the model's function within the MPC loop. The same type of tests is applied to the ADM1 predictors, and some results are displayed in Figure 7, showing equally good performance in this case.

3.3. Closed-loop performance of LSTM-MPC

The primary objective of the study was to evaluate the performance of the complete LSTM-MPC strategy. The controller was tasked with tracking step changes in the methane flow rate setpoint, while the plant emulators were subjected to the stochastic disturbances described in Section 3.1.

For the AM2 case (Figure 8), the controller achieves precise tracking with fast settling times after each setpoint change. The control action (dilution rate, D) exhibits rapid adjustments to counteract the system's faster dynamics and operate in a highly nonlinear region. Notably, the controller effectively uses the available control range, operating close to the defined constraints (dotted blue lines) to meet the performance objectives.

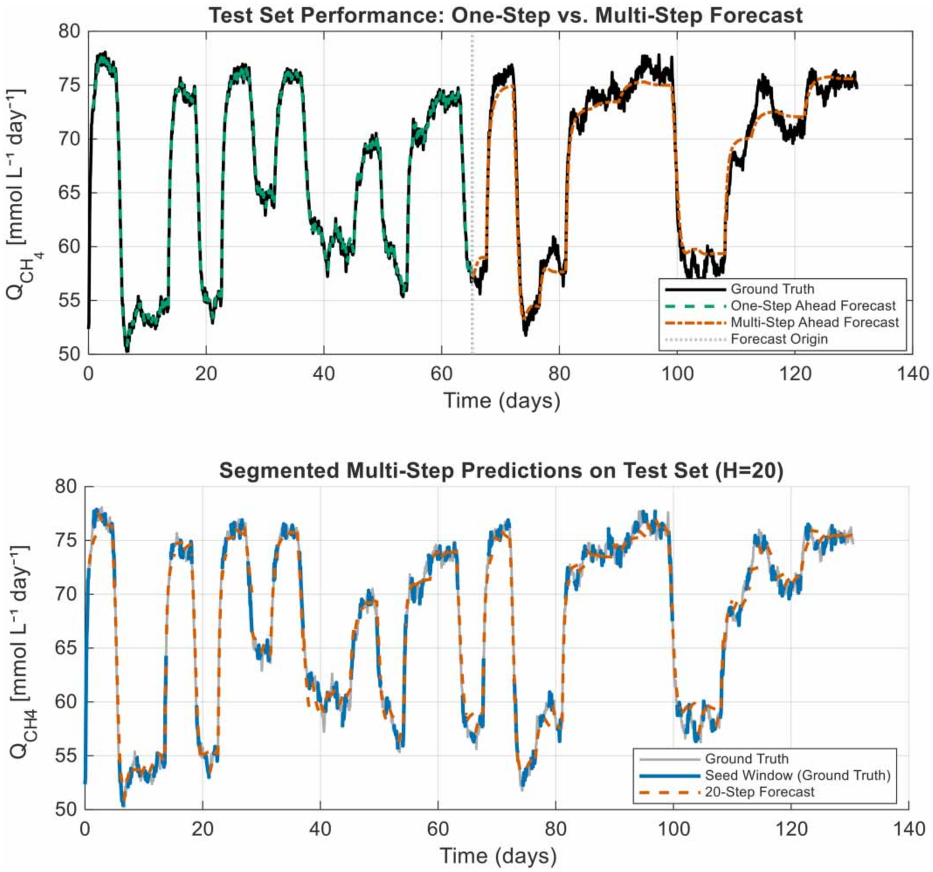

Figure 6 | Test performance showing combined one-step and multi-step predictions for the AM2 predictor. The bottom plot shows segmented multi-step predictions.

For the more complex ADM1 case (Figure 9), the controller also shows successful tracking, although the response is inherently slower, reflecting the process's sluggish dynamics. The control action (inlet flow rate, Q_{in}) is smoother and less aggressive compared with AM2, which is consistent with controlling a high-fidelity, slower system. Despite some minor oscillations around the setpoints, the controller maintains stability and effectively compensates for disturbances across a wide operational range.

3.4. Discussion

The results presented demonstrate the significant potential of combining data-driven LSTM predictors with an MPC structure for controlling AD processes. The main finding of the study is that, despite the inherent complexity of AD, which involves numerous interacting states and reactions, effective control can be achieved using a model trained on a minimal set of readily available online data: the methane flow rate and the manipulated dilution rate/inlet flow rate. This poses a significant practical advantage, as the online monitoring of other state variables, such as volatile fatty acid (VFA) or biomass concentrations, may correspond to a significant technical and economic challenge in industrial settings (Jamilis *et al.* 2018). The approach presented reduces the dependency on these complex measurements, relying instead on a variable that can commonly and reliably be measured in AD plants.

The successful performance of the LSTM-MPC against both a simplified (AM2) and a high-fidelity (ADM1) plant emulator, each subjected to stochastic disturbances, highlights the effectiveness of this data-driven approach. The LSTM network was capable of learning the input–output dynamics sufficiently well to allow the MPC to effectively counteract the effects of unmeasured disturbances and track setpoints. This suggests that detailed mechanistic models, while invaluable for process understanding, may not be strictly necessary for the controller design itself, provided that data of sufficient quantity and quality is available to train a reliable black-box predictor.

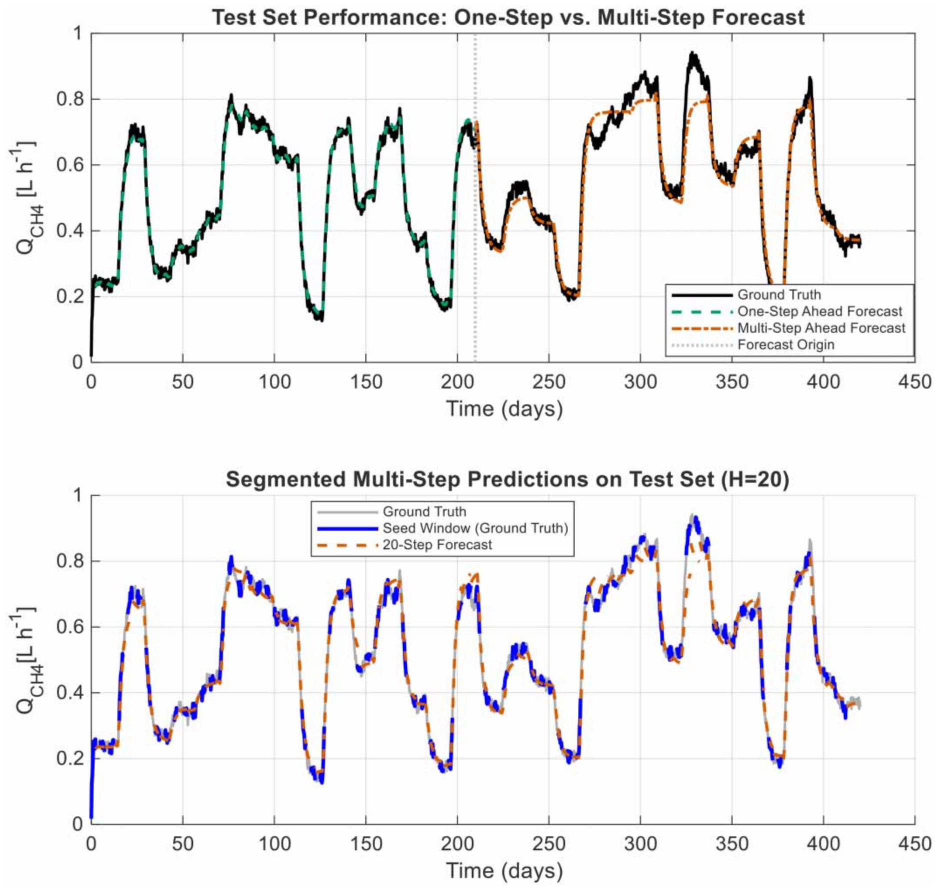

Figure 7 | Test performance showing combined one-step and multi-step predictions for the ADM1 predictors. The bottom plots show segmented multi-step predictions.

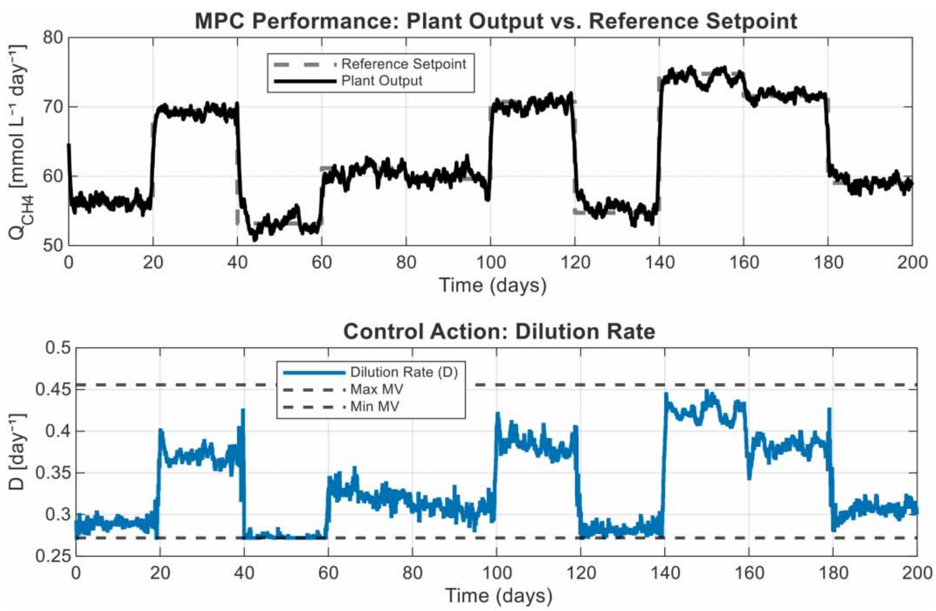

Figure 8 | Closed-loop performance of the LSTM-MPC showing setpoint tracking and control action with the AM2 emulator.

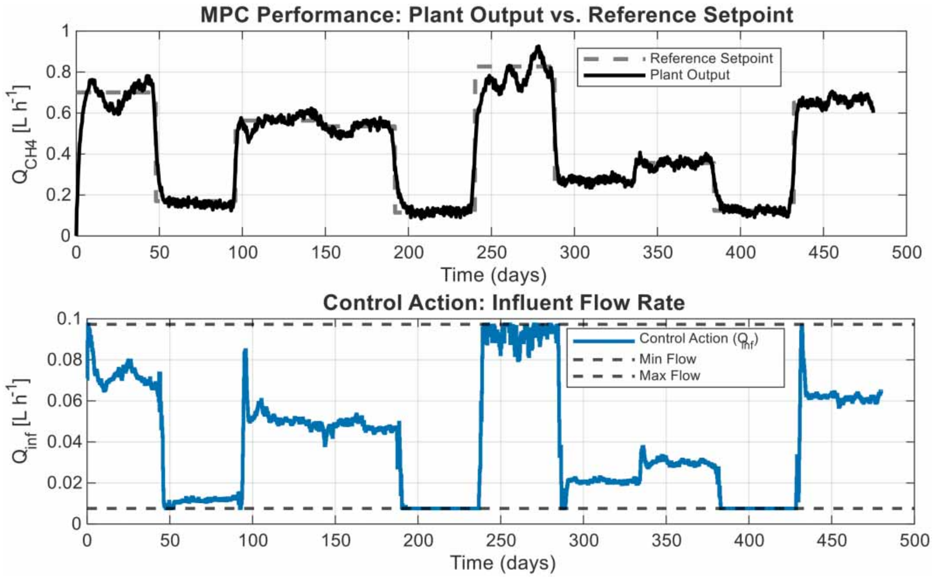

Figure 9 | Closed-loop performance of the LSTM-MPC showing setpoint tracking and control action with the ADM1 emulator.

3.4.1. Effect of data quantity

An additional analysis was performed to understand the impact of the amount of training data on the predictor's accuracy, specifically observing how the median RMSE for the methane flow rate changes with the number of perturbation steps in the training set.

As shown in Figure 10, increasing the number of perturbation steps initially drastically improves multi-step-ahead prediction accuracy. The trends of Figure 10 show that RMSE decreases for up to 5–6 perturbation steps before plateauing. Adding data beyond a certain point does not significantly improve performance.

The relative accuracy of the LSTM predictor is generally better for AM2 than ADM1, a difference likely attributable to the complexity of the underlying first-principle simulator.

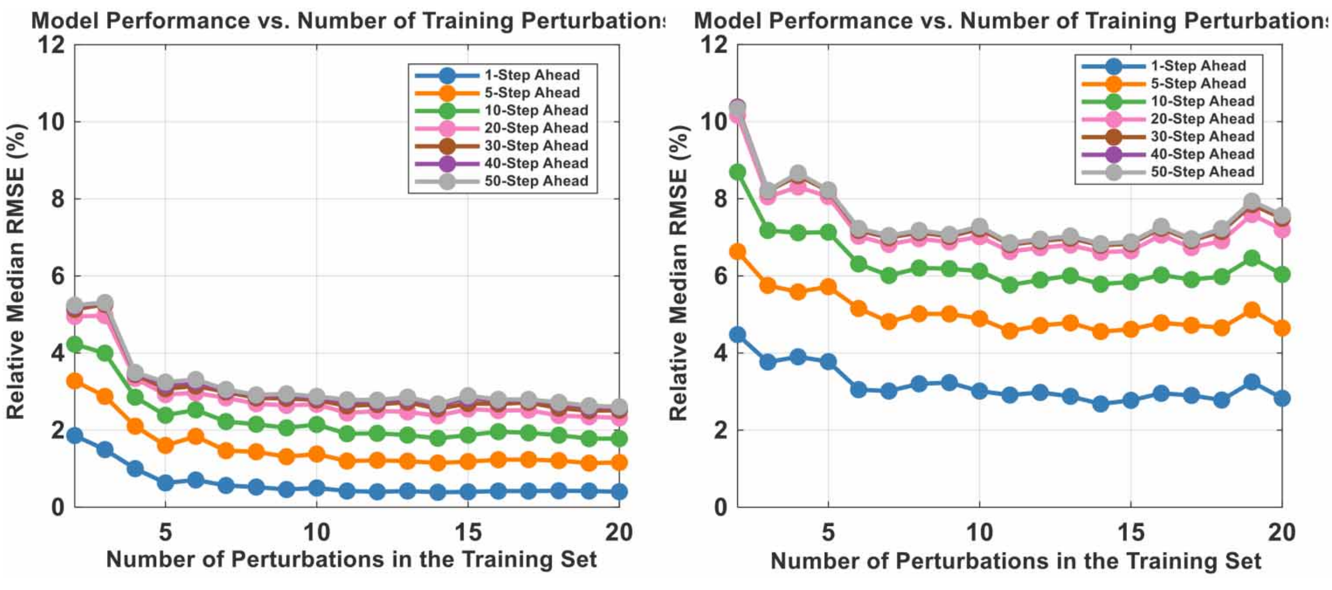

Figure 10 | Effect of training data on the *n*-step-ahead prediction of the LSTM model trained with the AM2 emulator (left) and with the ADM1 emulator (right). The plots show the median of 10 experiments vs. the number of step perturbations in the training set.

3.4.2. Computational feasibility and comparison with alternative strategies

Another critical aspect for real-world applicability is the computational demand of the controller. While complex neural networks such as LSTMs can be computationally intensive, the slow dynamics of the AD process provide a large time window between control actions. The simulations for this study were performed on a high-performance computer (AMD Ryzen 5 5600H, 32 GB RAM, and RTX 3060 Laptop); however, the results are indicative of real-time feasibility. As shown in Figure 11, the average computation time per MPC step was 1.60 s for AM2 and 4.38 s for ADM1. It should be noted that this includes the time to simulate one step of the plant emulator. In this configuration, the time margin is therefore several orders of magnitude smaller than the process sampling interval (e.g., ~1.75 h for AM2 and 6 h for ADM1). This suggests that the proposed LSTM-MPC approach may be computationally feasible and sufficient for real-time implementation.

To investigate the role of the model architecture, several alternative predictors were trained and evaluated in conjunction with the LSTM network using synthetic data from the ADM1 model. These included a gated recurrent unit (GRU) network, a simple feed-forward network based on a self-attention mechanism (ATTN), a standard feed-forward neural network in a NARX configuration (NARX), and a linear AutoRegressive with eXogenous inputs model (ARX). It is important to note that all models were trained using the same data and feedback structure, meaning that their predictions are based on a history of past process inputs and past process outputs. Additionally, a standard proportional integral (PI) controller was also evaluated.

Table 1 displays the results in normalized form (to make them CPU independent). The linear ARX model yields the poorest prediction and control performance among the MPC variants due to its inability to accurately represent the system's nonlinearity. The performance of the PI controller is even poorer, but is remarkable in the sense that it does not require an explicit model and is relatively robust. However, it rests on a careful tuning procedure, and is limited as it does not allow the

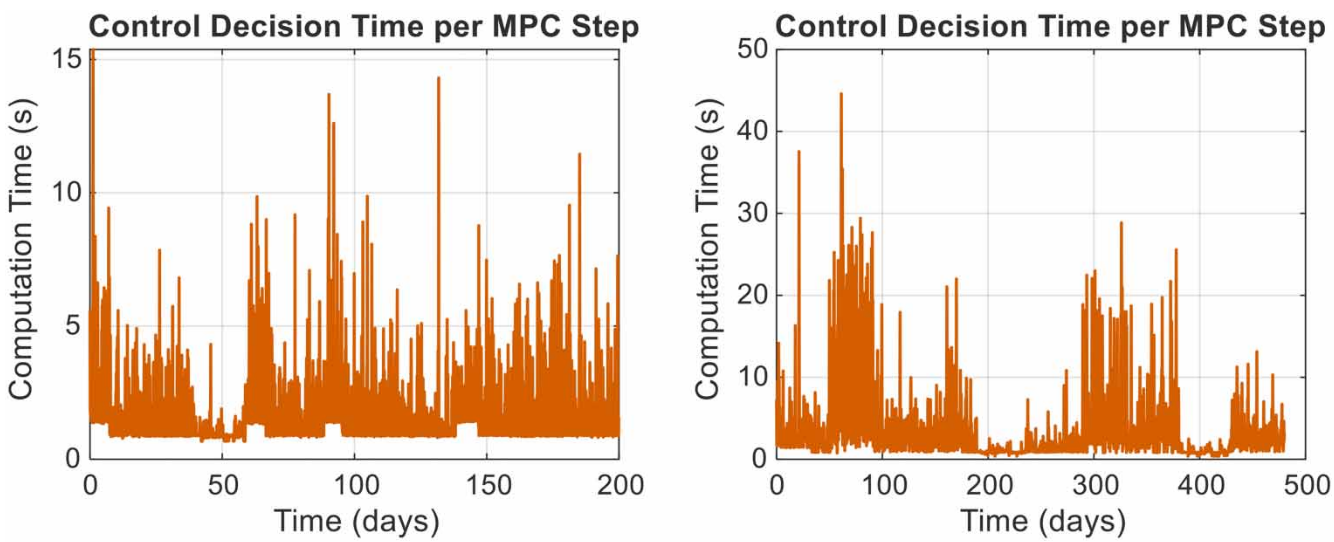

Figure 11 | Computation time per control step for the LSTM-MPC strategy applied to the AM2 (left) and ADM1 (right) emulators.

Table 1 | Performance and computational metrics for all evaluated controllers applied to the ADM1

Controller	RMSE	MAE	Mean time	Median time
LSTM-MPC	1	1	1	1
GRU-MPC	1.03	1.04	0.95	0.91
ATTN-MPC	1.10	1.16	0.37	0.52
NARX-MPC	1.54	1.43	0.41	0.53
ARX-MPC	1.18	1.36	3.0×10^{-4}	5.4×10^{-4}
PI	1.91	1.61	3.8×10^{-7}	5.7×10^{-7}

MAE, maximum absolute error

consideration of constraints on the input and output variables. MPC based on a nonlinear NN predictor outperforms PI but requires more computational resources. LSTM-MPC and GRU-MPC produce very similar performance. ATTN-MPC is an outsider with interestingly less computational demand.

3.4.3. Limitations and future research

While the results are promising, it is important to acknowledge the limitations of this simulation-based study, which are consistent with broader challenges in the application of ML to AD. Recent reviews have noted that model scalability, the scarcity of training data, and the poor interpretability of neural network models due to their 'black box' nature are limiting factors. As LSTMs are a type of neural network, they inherit this limitation, which hampers their adoption in industrial settings (Rutland *et al.* 2023).

In the present study, the predictors were trained on data from a single type of substrate. Sudden or significant changes in feedstock composition would likely lead to a model-plant mismatch, degrading controller performance. Moreover, while industrial processes are subject to complex, non-stationary disturbances (Lima Li & Appleby 2025), the stochastic noise introduced in this study via an AR(1) process represents only a simplified version of this variability. Finally, a significant challenge for data-driven approaches is the acquisition of informative training data. To generate this data, large, stepwise perturbations were applied to the manipulated variable in order to capture its corresponding step responses. Applying such deliberate and potentially destabilizing disturbances is often not feasible in a full-scale industrial plant, creating a practical bottleneck for training a model from scratch on a real system.

Investigating methods for online adaptation of the LSTM model would be a valuable next step to address the controller's robustness against long-term process changes. To mitigate the previously mentioned data acquisition bottleneck, transfer learning presents a promising strategy worth investigating (Weiss *et al.* 2016). This aligns with the call for more cross-domain knowledge transfer in the field to overcome data limitations (Rutland *et al.* 2023). The hypothesis is that an LSTM model could be pre-trained on a high-fidelity simulator, such as ADM1, and then be fine-tuned with a smaller, more accessible dataset from a real plant, potentially reducing the data collection burden for industrial implementation.

Future research should focus on validating this approach with synthetic data related to more complex AD processes, such as co-digestion of different substrates under various perturbations, and eventually, with experimental data from laboratory-scale or pilot reactors. As a first step, the authors are currently investigating the transferability of the approach to more complex anaerobic co-digestion processes as described in García-Gen *et al.* (2014).

Finally, exploring the inclusion of other easily measurable variables, such as pH and temperature, could also be studied to assess their potential impact on the prediction accuracy and robustness. It is likely that considering pH measurements could robustify the prediction in the event of reactor acidification, allowing the controller to take preventive actions and avoid harmful effects on the biomass, thereby preventing a subsequent reduction in methane flow.

4. CONCLUSIONS

The present study successfully demonstrated the feasibility and effectiveness of an MPC strategy using a data-driven LSTM network as the internal predictor for controlling complex AD processes. The key findings are as follows:

- LSTM networks proved to be highly effective predictors, capable of accurately capturing the nonlinear dynamics of both a simplified (AM2) and a high-fidelity (ADM1) AD process model using only commonly measured input–output data. The trained models showed stable open-loop prediction capabilities, a critical requirement for their use in an MPC framework.
- The integrated LSTM-MPC system achieved effective closed-loop performance. It successfully tracked methane flow rate setpoints and rejected unmeasured stochastic disturbances in both simulation environments, all while respecting operational constraints.
- The computational time required for online optimization was significantly lower than the chosen sampling interval, confirming the feasibility of this approach for real-time applications in slow-moving AD systems.

The primary contribution of this work is the demonstration that a purely data-driven, black-box control strategy can effectively manage AD processes. This offers a promising alternative to traditional model-based control, which often relies on complex, difficult-to-calibrate mechanistic models and requires measurements that are not readily available online.

While the simulation results are promising, future work must focus on experimental validation to ascertain real-world viability. Key research directions include the development of online adaptation mechanisms to enhance robustness, the investigation of transfer learning to mitigate data collection challenges for industrial implementation, and the evaluation of the strategy's scalability by extending it to anaerobic co-digestion and by assessing the impact of incorporating additional process variables.

ACKNOWLEDGEMENTS

This work was funded by ANID (Agencia Nacional de Investigacion y Desarrollo, Chile) through the project FONDECYT No. 11220818 granted to Santiago Garcia-Gen, and by Wallonie-Bruxelles International (WBI) through the 2023–2025 Projects of the VII Chile/Wallonia-Brussels Joint Commission.

DATA AVAILABILITY STATEMENT

All relevant data are available from an online repository or repositories: https://doi.org/10.6084/m9.figshare.29998807.

CONFLICT OF INTEREST

The authors declare there is no conflict.

REFERENCES

Batstone D. J., Keller J., Angelidaki I., Kalyuzhnyi S. V. , Pavlostathis S. G., Rozzi A., Sanders W. T. M., Siegrist H. & Vavilin V. A. (2002) The IWA Anaerobic Digestion Model No 1 (ADM1), *Water Sci Technol* **45** (10), 65–73. https://doi.org/10.2166/wst.2002.0292.

Beltramo, T. & Hitzmann, B. (2019) Evaluation of the linear and non-linear prediction models optimized with metaheuristics: application to anaerobic digestion processes, *Engineering in Agriculture, Environment and Food*, **12** (4), 397–403. https://doi.org/10.1016/j.eaef.2019.06.001.

Bernard, O., Hadj-Sadok, Z., Dochain, D., Genovesi, A. & Steyer, J. P. (2001) Dynamical model development and parameter identification for an anaerobic wastewater treatment process, *Biotechnology and Bioengineering*, **75** (4), 424–438. https://doi.org/10.1002/bit.10036.

Bolmanis, E., Dubencovs, K., Suleiko, A. & Vanags, J. (2023) Model predictive control – a stand out among competitors for fed-batch fermentation improvement, *Fermentation*, **9** (3), 206. https://doi.org/10.3390/fermentation9030206.

Dewasme, L. (2020) Neural network-based software sensors for the estimation of key components in brewery wastewater anaerobic digester: an experimental validation, *Water Science and Technology*, **80** (10), 1975–1985. https://doi.org/10.2166/wst.2020.019.

Dewasme, L., Sbarciog, M., Rocha-Cozatl, E., Haugen, F. & Vande Wouwer, A. (2019) State and unknown input estimation of an anaerobic digestion reactor with experimental validation, *Control Engineering Practice*, **85**, 280–289. https://doi.org/10.1016/j.conengprac.2019.02.003.

Donoso-Bravo, A., Mailier, J., Martin, C., Rodríguez, J., Aceves-Lara, C. A. & Vande Wouwer, A. (2011) Model selection, identification and validation in anaerobic digestion: a review, *Water Research*, **45** (17), 5347–5364. https://doi.org/10.1016/j.watres.2011.08.059.

Gaida, D., Wolf, C., Meyer, C., Stuhlsatz, A., Lippel, J., Bäck, T., Bongards, M. & McLoone, S. (2012) State estimation for anaerobic digesters using the ADM1, *Water Science and Technology*, **66** (5), 1088–1095. https://doi.org/10.2166/wst.2012.286.

García-Gen, S., Rodríguez, J. & Lema, J. M. (2014) Optimisation of substrate blends in anaerobic co-digestion using adaptive linear programming, *Bioresource Technology*, **173**, 159–167. https://doi.org/10.1016/j.biortech.2014.09.089.

García-Gen, S., Santos, L. O. & Vande Wouwer, A. (2022) Application of a nonlinear model predictive controller to the anaerobic digestion of readily biodegradable wastes, *IFAC-PapersOnLine*, **55**, 909–914. https://doi.org/10.1016/j.ifacol.2022.07.560.

Giovannini, G., Sbarciog, M., Steyer, J. P., Chamy, R. & Vande Wouwer, A. (2018) On the derivation of a simple dynamic model of anaerobic digestion including the evolution of hydrogen, *Water Research*, **134**, 209–225. https://doi.org/10.1016/j.watres.2018.01.036.

Gupta, R., Zhang, L., Hou, J., Zhang, Z., Liu, H., You, S., Ok, Y. S. & Li, W. (2023) Review of explainable machine learning for anaerobic digestion, *Bioresource Technology*, **369**, 128468. https://doi.org/10.1016/j.biortech.2022.128468.

Haugen, F., Bakke, R. & Lie, B. (2014) State estimation and model-based control of a pilot anaerobic digestion reactor, *Journal of Control Science and Engineering*, **2014**, 1–19. https://doi.org/10.1155/2014/572621.

Hochreiter, S. & Schmidhuber, J. (1997) Long short-term memory, *Neural Computation*, **9** (8), 1735–1780. 10.1162/neco.1997.9.8.1735.

Jamilis, M., Nuñnez, S., Garelli, F. & De Battista, H. (2018) 'Detectability and observability analysis of the anaerobic digestion process', *2018 Argentine Conference on Automatic Control (AADECA)*. IEEE, pp. 1–6.

Jimenez, J., Latrille, E., Harmand, J., Robles, A., Ferrer, J., Gaida, D., Wolf, C., Mairet, F., Bernard, O., Alcaraz-Gonzalez, V., Mendez-Acosta, H., Zitomer, D., Totzke, D., Spanjers, H., Jacobi, F., Guwy, A., Dinsdale, R., Premier, G., Mazhegrane, S., Ruiz-Filippi, G., Seco, A., Ribeiro, T., Pauss, A. & Steyer, J.-P. (2015) Instrumentation and control of anaerobic digestion processes: a review and some research challenges, *Reviews in Environmental Science and Biotechnology*, **14** (4), 615–648. https://doi.org/10.1007/s11157-015-9382-6.

Kil, H., Li, D., Xi, Y. & Li, J. (2017) Model predictive control with on-line model identification for anaerobic digestion processes, *Biochemical Engineering Journal*, **128**, 63–75. https://doi.org/10.1016/j.bej.2017.08.004.

Lima, D., Li, L. & Appleby, G. (2025) Biogas production modelling based on a semi-continuous feeding operation in a municipal wastewater treatment plant, *Energies*, **18** (5), 1065. https://doi.org/10.3390/en18051065.

Liu, Y., Tian, W., Xie, J., Huang, W. & Xin, K. (2023) LSTM-based model-predictive control with rationality verification for bioreactors in wastewater treatment, *Water (Switzerland)*, **15** (9), 1779. https://doi.org/10.3390/w15091779.

Mauky, E., Weinrich, S., Nägele, H., Jacobi, H. F., Liebetrau, J. & Nelles, M. (2016) Model predictive control for demand-driven biogas production in full scale, *Chemical Engineering and Technology*, **39** (4), 652–664. https://doi.org/10.1002/ceat.201500412.

McCormick, M. & Villa, A. E. P. (2019) LSTM and 1-D convolutional neural networks for predictive monitoring of the anaerobic digestion process. In: Tetko, I., Kůrková, V., Karpov, P., Theis, F. (eds) *Artificial Neural Networks and Machine Learning – ICANN 2019: Workshop and Special Sessions*. ICANN 2019. Lecture Notes in Computer Science, vol 11731. Cham, Switzertland: Springer. https://doi.org/10.1007/978-3-030-30493-5_65.

Rueda-Escobedo, J. G., Sbarciog, M., Moreno, J. A., Van Impe, J. & Vande Wouwer, A. (2022) Robust state and input estimation with enhanced convergence rate for monitoring anaerobic digestion, *Journal of Process Control*, **117**, 169–180. https://doi.org/10.1016/j.jprocont.2022.07.014.

Rutland, H., You, J., Liu, H., Bull, L. & Reynolds, D. (2023) A systematic review of machine-learning solutions in anaerobic digestion, *Bioengineering*, **10** (12), 1410. https://doi.org/10.3390/bioengineering10121410.

Salamattalab, M. M., Hasani Zonoozi, M. & Molavi-Arabshahi, M. (2024) Innovative approach for predicting biogas production from large-scale anaerobic digester using long-short term memory (LSTM) coupled with genetic algorithm (GA), *Waste Management*, **175**, 30–41. https://doi.org/10.1016/j.wasman.2023.12.046.

Weiss, K., Khoshgoftaar, T. M. & Wang, D. D. (2016) A survey of transfer learning, *Journal of Big Data*, **3** (1), 9. https://doi.org/10.1186/s40537-016-0043-6.

Zarzycki, K. & Ławryńczuk, M. (2021) LSTM and GRU neural networks as models of dynamical processes used in predictive control: a comparison of models developed for two chemical reactors, *Sensors*, **21** (16), 5625. https://doi.org/10.3390/s21165625.

First received 30 June 2025; accepted in revised form 1 September 2025. Available online 17 September 2025