

SURROGATE MODELING FOR THE OPTIMIZATION OF ADVANCED WATER TREATMENT SYSTEMS IN IRRIGATION AND WATER REUSE CONTEXTS

MODÉLISATION DE SUBSTITUTION POUR L'OPTIMISATION DES SYSTÈMES AVANCÉS DE TRAITEMENT DE L'EAU DANS LES CONTEXTES DE L'IRRIGATION ET DE LA RÉUTILISATION DE L'EAU

Derviş Durmaz¹, Rémy Dutto², Alain Rapaport³, Aymen Chaaben⁴, Jérôme Harmand²

ABSTRACT

This work develops a dynamic surrogate-modeling framework that approximates the input–output behavior of detailed mechanistic Anaerobic Membrane Bioreactor (AnMBR) models in view of its optimization for water reuse. The proposed approach uses a 30-day, one-minute-resolution dataset generated from the coupled biological–fouling model of Benyahia et al. (2024). Five neural-network architectures are benchmarked: a Long Short-Term Memory (LSTM) network, a Transformer, two recent time-series models (TimeMixer, TimeMachine), and a Stiff Physics-Informed Neural Network (Stiff-PINN). Within the training distribution, all five surrogates reproduce permeate flux with mean absolute percentage errors below 0.01 % and coefficients of determination above 0.96, with the Stiff-PINN reaching 0.0002 %. An explicit initial-condition sensitivity analysis, however, reveals a sharp dichotomy: when initial states fall outside the training distribution (pre-fouled membrane, combined biological–fouling stress), the four data-driven sequence models diverge by 290 % to 620 % MAPE, whereas the Stiff-PINN, which embeds the underlying ordinary differential equations, retains accuracy below 0.001%. The framework therefore identifies physics-informed surrogates as the architecture class best suited to support trustworthy optimization of membrane-based water reuse systems beyond historical operating regimes, while delivering the differentiability and batched evaluation required by gradient-based model-predictive control of irrigation-oriented reuse.

Keywords : Dynamic surrogate-modeling framework; Neural-network; Time-series models; Irrigation-oriented reuse.

RÉSUMÉ

Ce travail développe un cadre dynamique de modélisation de substitution qui reproduit approximativement le comportement entrée-sortie de modèles mécanistes détaillés de bioréacteur anaérobie à membranes (AnMBR), en vue d'optimiser la réutilisation de l'eau. L'approche proposée utilise un jeu de données couvrant 30 jours avec une résolution d'une minute, généré à partir du modèle couplé biologique-colmatage de Benyahia et al. (2024). Cinq architectures de réseaux neuronaux sont comparées : un réseau à mémoire à long et court terme (LSTM), un Transformer, deux modèles récents de séries temporelles (TimeMixer et TimeMachine) et un réseau neuronal informé par la physique pour systèmes raides (Stiff-PINN). Dans la distribution d'entraînement, les cinq modèles de substitution reproduisent le flux de perméat avec des erreurs absolues moyennes en pourcentage inférieures à 0,01 % et des coefficients de détermination supérieurs à 0,96, le Stiff-PINN atteignant 0,0002 %. Toutefois, une analyse explicite de la sensibilité aux conditions initiales révèle une nette dichotomie : lorsque les états initiaux se situent hors de la distribution d'entraînement — membrane précolmatée ou contraintes biologiques et de colmatage combinées —, les quatre modèles séquentiels fondés sur les données divergent, avec des valeurs de MAPE comprises entre 290 % et 620 %, tandis que le Stiff-PINN, qui intègre les équations différentielles ordinaires sous-jacentes,

1 Artificial Intelligence and Agricultural Technologies Research Center, TAGEM, Ankara, Türkiye

2 LBE, Univ. Montpellier, INRAE, Narbonne, France

3 UMR MISTEA, Univ. Montpellier, INRAE, Montpellier, France

4 IEM, Univ. Montpellier, Montpellier, France

conserve une erreur inférieure à 0,001 %. Le cadre désigne ainsi les modèles de substitution informés par la physique comme la classe d'architectures la mieux adaptée à une optimisation fiable des systèmes de réutilisation de l'eau à membranes au-delà des régimes de fonctionnement historiques, tout en offrant la différentiabilité et l'évaluation par lots nécessaires à la commande prédictive par modèle fondée sur le gradient pour la réutilisation de l'eau destinée à l'irrigation.

Mots clés : Cadre dynamique de modélisation de substitution; Réseaux neuronaux; Modèles de séries temporelles; Réutilisation de l'eau destinée à l'irrigation.

Introduction

The intensification of water scarcity and the growing demand for reliable water supplies have led to increased interest in high-performance water treatment technologies, including advanced membrane-based systems. These technologies offer high treatment efficiency and water quality, making them particularly attractive for applications related to water reuse, especially for irrigation purposes, and integrated water resources management. However, their widespread implementation is often limited by high investment and operational costs, as well as by their sensitivity to operating conditions. Optimizing the operation of such systems is therefore a key challenge to improve their economic and environmental performance while ensuring robustness and reliability, (Sandoval-Garcia et al., 2025).

The development and evaluation of advanced optimization and control strategies traditionally rely on experimental pilot plants or full-scale facilities. While these platforms provide valuable insights, they are costly to operate, require long experimental campaigns, and are generally restricted to a narrow range of operating conditions. These constraints limit the systematic exploration of system behavior, particularly under dynamic or extreme scenarios that are relevant for irrigation-oriented water reuse applications, such as fluctuating demands, variable influent quality, or intermittent operation, (Chaaben et al., 2025).

An alternative approach consists in the use of virtual systems based on high-fidelity process models to assess and validate optimization methodologies. For membrane-based water treatment systems, such mechanistic models can accurately describe transport processes, fouling dynamics, and operational constraints (Sandoval et al., 2025). However, their high level of detail comes at the expense of significant computational burden. Simulation may lose accuracy, and times can become prohibitive, especially when these models are embedded in optimization or decision-support frameworks that require repeated evaluations over long time horizons, (Quo and Wang, 2008).

To address this challenge, this work proposes a surrogate modeling framework aimed at accelerating the optimization of high-performance water treatment systems. Increasingly used in Artificial Intelligence, dynamic neural-network-based surrogate models are developed to approximate the input–output behavior of detailed mechanistic models (Frontistis et al., 2023; Abuwatfa et al., 2023). Recurrent architectures, such as Long Short-Term Memory (LSTM) networks or Recurrent Neural Networks (RNNs), are particularly suitable for capturing the nonlinear and time-dependent dynamics characteristic of water treatment processes. Once trained, these black-box models provide near-instantaneous simulations, making them well adapted for use in optimization, real-time control, and scenario analysis.

The proposed approach follows a structured three-phase methodology. In the first phase, comprehensive datasets are generated using advanced mechanistic models under a wide range of operating conditions representative of a water treatment and reuse framework for irrigation. Special attention is paid to the selection of initial conditions, as they constitute one of the main sources of uncertainty in the simulation of physically accurate models, particularly for systems exhibiting slow dynamics or memory effects such as fouling. In the second phase, surrogate models are trained to learn the dynamic mapping from control inputs to system outputs, with explicit analysis of their sensitivity to initial conditions. This sensitivity analysis is essential to assess the robustness of the surrogate models and their ability to reproduce transient behavior under realistic operating scenarios.

In the third phase, the surrogate models are validated by confronting their predictions with those of the high-fidelity models in operating conditions that were not included in the training dataset.

An approach based on surrogate modelling

The framework adopted in this work approximates a high-fidelity (here a mechanistic) dynamical model of the form $\dot{x}(t)=f(x(t),u(t))$, $y(t)=g(x(t),u(t))$ by a neural surrogate that maps control inputs and an initial state to the trajectory of the output of interest. The mechanistic simulator is used solely as a virtual experimental rig: it generates training trajectories under a diverse range of operating conditions, including the abrupt regime changes that characterize membrane filtration cycles. Once trained, the surrogate is used in place of the simulator within optimization, control, and scenario-analysis time-limited loops, where its low marginal cost per evaluation and its end-to-end differentiability are decisive.

Two design choices govern the quality of the resulting surrogate. The first is the design of the data-generation protocol: the mechanistic model must be excited by inputs that span the operating envelope expected during deployment, and the initial conditions must cover not only nominal start-up but also realistic stress states (high biomass, pre-fouled membrane). The second is the choice of architecture, which determines whether the surrogate generalizes beyond the operating regimes seen during training. We compare both purely data-driven sequence models, which learn statistical regularities from the training trajectories, and a physics-informed network whose loss embeds the residual of the governing ordinary differential equations (ODEs). The latter inherits the structural inductive bias of the mechanistic model and is expected to remain accurate in extrapolative regimes.

Application: predicting fouling of a decentralized AnMBR for water reuse

In order to evaluate the approach, we consider a model initially proposed in (Benyahia et al., 2024). This model couples a modified dynamic Anaerobic Digestion (AD) model, initially proposed in (Bernard et al., 2001), to which SMP dynamics was added, with a membrane fouling model involving two resistances, namely the cake and the pore blocking resistances. The last one is considered to be mainly due to the presence of SMP. The main assumptions of the model are as follows:

- The system is supposed to operate at constant TransMembranePressure (TMP).
- The accumulation of particulate and part of the soluble matter (notably the SMP) are acting on the cake layer and on the internal resistance due to pore blocking in such a way, over time, both resistances are increasing and the available filtration surface decreases, further decreasing the flow rate able to pass through the membrane. This is illustrated in Figure 1.
- The biological states are X_1 , X_2 , S_1 and S_2 are the acidogenic and methanogenic bacteria, the COD and the VFA concentrations, respectively. In addition, SMP dynamics was also included.




Figure 1. Membrane fouling by cake formation and pores clogging, from (Benyahia et al., 2024)

The system is operating either in filtration or in backwash mode. We do not recall all dynamical equations here since they can be found in (Benyahia et al., 2024). However, we recall the main principles which support the modeling of the filtration part since it is critical for our objectives. The flux passing through the membrane is given by J :

$$J(t) = \frac{Q_{out}}{A(t)} = \frac{\Delta P}{\mu(R_0 + R(t))} \quad (1)$$

where Q_{out} is the flow rate passing through the membrane, $A(t)$ is the available surface through which the filtrate can pass, ΔP is the TMP, μ is the permeate viscosity, R_0 is the initial membrane resistance and $R(t)$ is the resistance of the membrane. Using a Resistance In Series (RIS) model, the resistance is supposed to be modeled as the sum of the cake and pore blocking resistances as:

$$R(t) = R_c(t) + R_p(t) \quad (2)$$

Introducing m_c and m_p , respectively the mass of matter accumulated onto the membrane and into its pores, it is possible to relate the state variables of the biological model to the state variables describing the filtration dynamics during filtration and backwash.

The numerical values used in the simulations follow Benyahia et al. (2024, Table A1): $\Delta P = 1.5$ bar, $\mu = 10^{-3}$ Pa.s, $A_0 = 1$ m², and $R_0 = 1.11 \times 10^{13}$ m⁻¹ (the intrinsic membrane resistance derived from a clean-water test). With these parameters and a clean membrane at $t = 0$, Eq. (1) gives an initial permeate flux $J(0) = \Delta P / (\mu R_0) \approx 1.35 \times 10^{-5}$ m s⁻¹, which corresponds to ≈ 48.6 L m⁻² h⁻¹ — a typical clean-membrane operating point for a tubular AnMBR at 1.5 bar TMP. This baseline reflects the clean-membrane parameter set of Benyahia et al. (2024, Table A1) rather than the partially fouled conditions of their experimental calibration (Section 4 of Benyahia et al., 2024); the qualitative dynamics of the coupled biological–fouling system are unaffected by this absolute scaling.

Objectives and surrogate structure

Biological models exhibit very different time constants (from seconds for physico-chemical equilibria to days or months for biomass dynamics), and their coupling to filtration modules whose operation alternates between filtration and backwash adds further stiffness. This combination makes direct integration of the equations time-consuming and is an obstacle to real-time use. The objective is therefore to use surrogate modeling to predict rapidly the variables of interest as functions of the main input variables. The structure of the surrogate model involves 7 state variables (X_1 , X_2 , S_1 , S_2 , SMP, m_c and m_p), 3 control inputs (Q_{in} , S_{1in} and S_{2in}) that are the input flow rate, the input COD and the input VFA concentrations, respectively and a key output consisted in the Permeate flux $J(t)$, cf. Figure 2. It also accounts for backwash/relaxation phases through the use of a binary control signal.

Because the AnMBR model couples slow biological dynamics with fast filtration dynamics, the surrogate is designed to predict the state vector $x = (X_1, X_2, S_1, S_2, \text{SMP}, m_c, m_p)$, and the permeate flux $J(t)$ is then recovered from the predicted state through the Darcy relation given in equation (1). Two contrasted training paradigms are considered, illustrated in Figure 2. The four sequence models (LSTM, Transformer, TimeMixer, TimeMachine) are trained as window-to-window predictors: each sample is a 24-minute window of past states and past controls, augmented with the future control trajectory, and the network outputs the next 24 minutes of states. At deployment they therefore require a short warm-up window of states, after which they are rolled forward autoregressively along any prescribed control trajectory. The Stiff-PINN follows a different paradigm: at training time it uses a single time point $(x(t), u(t))$ and is asked to output the state derivative dx/dt , with a residual loss that penalizes its deviation from the analytical right-hand side of the AnMBR ODEs (Benyahia et al., 2024). At deployment it is therefore integrated forward from any chosen initial condition $x(0)$ along the future control trajectory $u(t..T)$ without a warm-up window — exactly the form required by gradient-based model-predictive control.




Figure 2. Input–output structure of the two surrogate paradigms. The sequence models (LSTM, Transformer, TimeMixer, TimeMachine) learn a 24-minute window-to-window mapping and are rolled out autoregressively after a warm-up window. The Stiff-PINN learns the vector field $dx/dt = NN(x, u)$ at single time points and is integrated forward from any initial condition $x(0)$. For both paradigms, the permeate flux $J(t)$ is recovered from the predicted state through the Darcy relation given by equation (1).

A dataset was generated by integrating the Benyahia et al. (2024) model over 30 days of operation at one-minute resolution (43,200 samples), with stochastic time-varying inputs and periodic 45-minute filtration/15-minute relaxation cycling. The dataset was split chronologically into 70 % training, 15 % validation and 15 % test (cf. Figure 3).

Five neural-network architectures were trained on the same windows and loss function: a Long Short-Term Memory (LSTM) network as a sequence baseline; an encoder-only Transformer; TimeMixer (Wang et al., 2024), a multi-layer-perceptron model with past-decomposable mixing; TimeMachine (Ahamed and Cheng, 2024), a selective state-space (Mamba-type) model; and a Stiff Physics-Informed Neural Network (Stiff-PINN), a compact multilayer perceptron whose loss penalizes the residual of the governing ODEs.




Figure 3. Operational dataset generated from the Benyahia et al. (2024) mechanistic model. Left column: 30-day evolution of permeate flux J , fouling masses m_c and m_p , and time-varying control inputs Q_{in} and $S_{1,in}$. Right column: 4-hour zoom highlighting the filtration/relaxation cycle (grey bands).

Results

In-distribution accuracy

Table 1 reports the predictive performance of the five surrogates on the held-out chronological test set, together with their inference and training times on a single GPU. All five architectures reproduce the permeate flux with mean absolute percentage errors (MAPE) below 0.01 % and coefficients of determination R^2 above 0.96, confirming that the AnMBR input–output mapping lies well within the representational capacity of modern sequence models. The Stiff-PINN attains the highest accuracy (MAPE = 0.0002 %, R^2 = 0.9955), followed by TimeMixer (0.0036 %), the LSTM baseline (0.0048 %), TimeMachine (0.0061 %) and the Transformer (0.0071 %). Figure 4 shows that, at this scale, the predicted flux of all surrogates is visually indistinguishable from that of the mechanistic model; the residual panel makes it explicit that the prediction errors remain bounded below ± 0.012 % over the entire 6-day filtration window.

It must be stressed that, for a single forward simulation over 10 days, the surrogates do not outperform the well-conditioned BDF integrator applied to the AnMBR model (≈ 0.04 s for the mechanistic ODE versus 0.08–0.15 s for the data-driven surrogates and 5.09 s for the step-by-step Stiff-PINN integrator). The advantage of the surrogate framework therefore lies elsewhere: in end-to-end differentiability with respect to control inputs and initial conditions, in batched evaluation of thousands of trajectories at near-constant marginal cost on standard GPU hardware, and in deployment in embedded control settings where the full ODE parameterization is unavailable. These properties are precisely those required by gradient-based model-predictive control and by uncertainty quantification.

Table 1. Predictive accuracy and computational cost of the five surrogate architectures on the in-distribution test set.

Model	MAPE (%)	R ²	Inference (s)	Training (s)	Parameters
Stiff-PINN	0.0002	0.9955	5.09	142	~9.5k
TimeMixer	0.0036	0.9912	0.13	12	~92k
LSTM	0.0048	0.9826	0.09	11	~84k
TimeMachine	0.0061	0.9684	0.08	8	~110k
Transformer	0.0071	0.9663	0.15	22	~136k




Figure 4. Prediction of permeate flux on the held-out test set. Top: filtration-phase flux trajectories of the five surrogates (coloured) overlaid on the mechanistic ground truth (black) over a 6-day window. Bottom: relative prediction errors stay bounded within ± 0.012 % for all architectures.

Sensitivity to initial conditions

The most informative result emerges from the explicit initial-condition (IC) sensitivity analysis. Six initial conditions were defined: four within the operating envelope sampled during training (nominal, high biomass, high substrate load, low biomass) and two explicitly extrapolative (a pre-fouled membrane with $m_c = 2.0$ g, two orders of magnitude beyond the training maximum, and a combined biological-fouling stress state). For each IC, both the mechanistic model and the trained surrogates were integrated over 10 days, and the resulting flux trajectories are compared in Figure 5.




Figure 5. Initial-condition sensitivity of the five surrogates. Panels (a)–(d): in-distribution ICs, all surrogates indistinguishable from the mechanistic ground truth. Panels (e)–(f), highlighted: extrapolative ICs (pre-fouled membrane and combined biological-fouling stress); the four data-driven sequence models diverge by several hundred percent and revert to the in-distribution mean flux, while the Stiff-PINN (red dashed) tracks the mechanistic trajectory.

For the four in-distribution initial conditions, the five architectures all reproduce the mechanistic flux trajectory with MAPE below 0.01 %. For the two extrapolative initial conditions, however, a sharp dichotomy appears. The four purely data-driven sequence models lose physical consistency: their predicted trajectories on the pre-fouled membrane and the combined biological-fouling stress state exhibit large transient oscillations of physically inadmissible magnitude during the first few hours and then settle far from the mechanistic reference, with cumulative MAPE ranging from 290 % to 620 %. This behaviour is consistent with surrogates that have learned the statistical regularities of the training trajectories but have not internalised the underlying physics: when initialised outside the training envelope, no statistical prior is available to reconstruct the correct dynamics. The Stiff-PINN, by contrast, remains accurate (MAPE below 0.001 %): the physics-residual term in its loss has forced the network to learn the vector field of the governing ODEs rather than only the trajectories sampled during training. Figure 6 quantifies this divergence on a logarithmic scale.




Figure 6. Quantitative summary. Left: in-distribution MAPE on the test set (linear scale); all surrogates below 0.01 %. Right: out-of-distribution MAPE on the pre-fouled membrane IC (logarithmic scale); the Stiff-PINN is the only architecture to retain accuracy (0.001 %), while sequence models lose four to six orders of magnitude.

This finding is the central contribution of the work. It shows that high statistical accuracy on an in-distribution test set is not sufficient to qualify a surrogate for use in optimization studies that explore extrapolative operating points: an optimizer that searches over feed concentrations, filtration cycles or membrane states will routinely visit operating points that are absent from the training set, and only physics-informed surrogates retain meaningful predictions in this regime. From a practical standpoint, two complementary recommendations follow. For tasks that remain within the operating envelope of historical data - model-predictive control of nominal cycles, fault detection - TimeMixer offers the best compromise between accuracy, training cost and inference speed. For tasks that explore extrapolative regimes - design optimization, robust scheduling under uncertain influent quality, what-if analyses on accelerated fouling - the Stiff-PINN is the only surrogate that retains physical consistency. A hybrid pipeline, in which a fast data-driven surrogate is locally retrained with PINN-generated trajectories whenever the optimizer leaves the training envelope, appears to be a promising avenue for future work.

Conclusions

In this paper, five neural-network surrogate modeling architectures such as LSTM, Transformer, TimeMixer, TimeMachine and a Stiff Physics-Informed Neural Network were benchmarked for the rapid estimation of permeate flux in an AnMBR used for water reuse. Among the important results, one may insist on the fact that (i) all five surrogates achieve sub-0.01% MAPE in-distribution, validating the surrogate-modeling approach for AnMBR fouling prediction; (ii) only physics-informed networks retain accuracy on extrapolative initial conditions, which is the relevant regime for

optimization; (iii) the framework is directly applicable to the question of non-conventional water reuse for irrigation, where AnMBR units are candidates for the polishing of secondary effluents in peri-urban agriculture. The results demonstrate that the proposed surrogate models significantly reduce computational time while maintaining high accuracy and robust dynamic behavior. This enables extensive testing and comparison of optimization strategies for water treatment systems under variable and uncertain conditions. By facilitating the rapid evaluation of operational strategies, the proposed framework supports the development of efficient, resilient, and cost-effective water treatment solutions, contributing to sustainable water management and irrigation systems. Future work will integrate the Stiff-PINN within a model-predictive controller for the optimization of permeation/relaxation cycles under stochastic feed quality following the approach initially proposed in (Chaaben et al., 2026).

Acknowledgements

This work has been done within the framework of the COST Action CA23104 Water4REUSE.

References

- Abuwatfa, W.H., N. AlSawaftah, N. Darwish, W.G. Pitt and G.A. Hussein (2023) A review on membrane fouling prediction using artificial neural networks (ANNs). *Membranes*, 13, 685, <https://doi.org/10.3390/membranes13070685>
- Ahamed, M.A., and Q. Cheng (2024) TimeMachine: A time series is worth 4 mambas for long-term forecasting. 27th European Conference on Artificial Intelligence (ECAI 2024), Santiago de Compostela, Spain.
- Benyahia, B., A. Charfi, G. Lesage, M. Heran, B. Cherki and J. Harmand (2024) Coupling a Simple and Generic Membrane Fouling Model with Biological Dynamics: Application to the Modeling of an Anaerobic Membrane BioReactor (AnMBR). *Membranes*, 14, 69. <https://doi.org/10.3390/membranes14030069>
- Bernard, O., Z. Hadj-Sadok, D. Dochain, A. Genovesi and J.P. Steyer (2001), Dynamical model development and parameter identification for an anaerobic wastewater treatment process. *Biotechnol. Bioeng.*, 75: 424-438. <https://doi.org/10.1002/bit.10036>
- Chaaben, A., F. Ellouze, N. Ben Amar, A. Rapaport, M. Heran and J. Harmand (2025) Evaluation of the genericity of an adaptive optimal control approach to optimize membrane filtration systems, *Membranes*, <https://doi.org/10.3390/membranes15060157>
- Frontistis, Z., G. Lykogiannis, A. Sarmpanis (2023) Machine learning implementation in membrane bioreactor systems: Progress, challenges, and future perspectives: A review. *Environments*, 10, 127, <https://doi.org/10.3390/environments10070127>
- Quo, C.F. and M.D. Wang (2008) Quantitative analysis of numerical solvers for oscillatory biomolecular system models. *BMC Bioinformatics* 9 (Suppl 6), S17. <https://doi.org/10.1186/1471-2105-9-S6-S17>
- Sandoval-Garcia, V., M.V. Ruano, M. Alliet, C. Brepols, J. Comas, J. Harmand, M. Heran, G. Mannina, I. Rodriguez-Roda, I. Smets and A. Robles (2025) Modeling MBR fouling: a critical review analysis towards establishing a framework for good modeling practices, *Water Res.*, 268, <https://doi.org/10.1016/j.watres.2024.122611>
- Wang, S., H. Wu, X. Shi, T. Hu, H. Luo, L. Ma, J.Y. Zhang, and J. Zhou (2024) TimeMixer: Decomposable multiscale mixing for time series forecasting. International Conference on Learning Representations (ICLR 2024), Vienna, Austria.