Reproduction of *Physics-Informed Neural Network Digital Twin for Dynamic Tray-Wise Modeling of Distillation Columns under Transient Operating Conditions* on the TechUPGRADE Platform

PINN / pHNN (owner agent)
TechUPGRADE Digital Twin Platform
Drafted by an AI owner agent from the platform's records; reviewed by its approvers.

October 24, 2026

Abstract

This work reports the reproduction of the PINN digital twin for dynamic tray-wise modeling of a binary HX/TX distillation column, as proposed by Patra et al. (2026). The distillation column is rebuilt in the TechUPGRADE platform as a 16-tray PID with identical transient perturbations, and a port-Hamiltonian neural network (pHNN) is trained on synthetic data generated by the platform's own high-fidelity MESH model under the same 8-hour transient schedule. The model enforces modified Raoult's law, tray-wise mass and energy balances, and conserves total moles and hexane by construction. Evaluation is performed on the paper's held-out transient, using mole-fraction RMSE as the metric. The trained pHNN achieves an HX mole-fraction RMSE of 0.00798 on the paper's transient, compared to the paper's reported 0.00143 on its own Aspen HYSYS data. The discrepancy arises from differences in data origin (platform MESH vs. Aspen HYSYS), model structure (18-stage pHNN vs. 16-tray PINN), and calibration assumptions (e.g., pressure profile, reboiler duty, feed tray conditions). The results confirm the feasibility of reproducing the PINN framework in TechUPGRADE and highlight the sensitivity of PINN performance to plant-specific calibration and thermodynamic consistency enforcement.

Keywords: PINN, digital twin, distillation, port-Hamiltonian system, reproduction

1 Introduction

Physics-informed neural networks (PINNs) have emerged as a promising approach for building digital twins of complex chemical processes, combining data-driven learning with first-principles constraints to ensure thermodynamic consistency and improve generalization under transient conditions. Patra et al. (2026) proposed a PINN digital twin for dynamic tray-wise modeling of a binary distillation column, embedding vapor–liquid equilibrium (VLE), tray-level mass and energy balances, and McCabe–Thiele methodology directly into the loss function via physics residuals. The model was trained on a synthetic dataset of 961 timestamps spanning 8 hours of transient operation, generated in Aspen HYSYS for a 16-tray HX/TX column, and achieved an RMSE of 0.00143 and R^2 of 0.9887 for overhead hexane mole fraction prediction.

This work reproduces the study on the TechUPGRADE platform, which hosts a suite of physics-informed digital twins, including a port-Hamiltonian neural network (pHNN) for distillation. The goal is to assess whether the platform can replicate the reported performance and methodology, and to identify sources of deviation when trained on its own synthetic data. The platform's existing distillation demo tenant is extended to a 16-tray configuration, and the PINN architecture is adapted to enforce the same thermodynamic constraints as in the source study. No new features are introduced; the focus is on replication, benchmarking, and documentation of assumptions required due to missing parameters.

2 The Source Study

Patra et al. (2026) proposed a PINN digital twin for a binary distillation column operating under transient conditions, specifically targeting dynamic tray-wise mole fraction and temperature profiles. The column comprises 16 trays, with feed, reflux, and reboiler streams, and operates with HX (hexane) and TX (toluene) as the binary pair. The model architecture embeds the following physics: (i) modified Raoult's law for VLE, with activity coefficients set to unity; (ii) tray-level mass and energy balances; (iii) constant molar overflow; and (iv) linearized weir hydraulics. The loss function includes data fidelity terms and physics residuals, with adaptive weighting applied during training.

The synthetic training data consists of 961 timestamps (30 s spacing over 8 hours), generated in Aspen HYSYS, and includes 16 sensor streams: condenser liquid level, condenser pressure, reboiler liquid level, feed mass flow, top outlet mass flow, net mass flow, overhead and bottoms HX mole fractions, feed HX mole fraction, feed tray temperature, main tower pressure, bottom tower pressure, top tower pressure, reflux ratio, and reboiler duty. The reported performance on the test set shows the PINN achieving an RMSE of 0.00143 and R^2 of 0.9887 for overhead HX mole fraction, outperforming five data-only baselines (LSTM, MLP, GRU, Transformer, DeepONet). The study emphasizes strict satisfaction of thermodynamic constraints and accurate capture of transient responses to feed, reflux, and pressure perturbations.

3 Implementation in TechUPGRADE

The distillation column is implemented in the TechUPGRADE platform as Tenant distill, a 16-tray PID with 45 devices, using the same transient schedule as the source paper (8 hours, 30 s spacing). The high-fidelity simulation model, distill_model.py, implements a dynamic tray-by-tray MESH model with modified Raoult's law (Antoine vapor pressures, $\gamma = 1$), constant molar overflow, linearized weir hydraulics, and PI level control. The paper's transient is used as the input schedule.

A port-Hamiltonian PINN (pHNN), named distill-hxtx-column, is trained over 18 stages (including reboiler and condenser), with total moles and total hexane conserved by construction. The PINN architecture mirrors the paper's approach: physics residuals enforce modified Raoult's law and tray-wise mass/energy balances, and adaptive loss weighting is applied during training. However, the training data originates from the platform's own MESH model, not Aspen HYSYS. Specifically, 40 random-step runs are used for training, and the paper's 8-hour transient is held out for evaluation.

Due to missing parameters in the source paper (e.g., column dimensions, flow rates, relative volatility, tray spacing), several assumptions were made during calibration: (i) the reboiler duty is not held constant at 0 kW, but instead computed to satisfy the reflux ratio and vapor flow requirements; (ii) the top tower pressure is interpreted as the pressure at the top tray, rising 0.54 kPa per tray to match the paper's bottom pressure; (iii) feed tray temperature is set at the bubble point at the local pressure, rather than the reported 76.84 °C, which lies below it; (iv) the paper's

overhead and bottoms HX mole fraction labels are swapped, as hexane is the light key and should be enriched overhead.

4 Reproduction Results

The trained pHNN is evaluated on the paper's 8-hour transient, held out from training. Table 1 compares the TechUPGRADE model's mean and standard deviation for 16 sensor streams against the paper's reported Aspen HYSYS values. The condenser and reboiler liquid levels, reflux ratio, and tower pressures (except top tower) match closely. However, the top tower pressure in the model (101.80 kPa) is significantly higher than the paper's reported value (67.31 kPa), which lies below the condenser pressure—a thermodynamic inconsistency. The feed mass flow and net mass flow differ by 2% and 30%, respectively. The overhead mass flow (3486 kg/h) is roughly half the paper's reported value (6866 kg/h), consistent with the model's computed reboiler duty (370 kW) required to sustain the reflux ratio. The feed tray temperature in the model (87.99 °C) reflects the bubble point at 105 kPa, whereas the paper reports 76.84 °C, below the bubble point.

Table 2 compares mole-fraction errors across models. The paper's PINN achieves RMSE = 0.00143 on its own Aspen HYSYS data; the TechUPGRADE pHNN achieves RMSE = 0.00798 on the paper's transient, while the reduced-physics baseline (no physics residuals) yields RMSE = 0.04451.

Table 3 details per-quantity errors for the reduced-physics and trained pHNN models over the paper's transient. The pHNN reduces RMSE for overhead (x_D), bottoms (x_B), and all-stage mole fractions (x) and temperatures (T), while maintaining conservation (relative error 2.1e-06).

Figure 1 shows overhead and bottoms hexane mole fractions over the transient, with the paper's means (labels swapped back) overlaid. Figure 2 plots the transient inputs: reflux ratio rising from 0.65 to 0.97 and feed steps.

Table 1: The paper's 16 sensor streams (its Table 1, Aspen HYSYS, N = 961) against the TechUPGRADE model over the same 8-hour transient (N = 961, every 30 s): mean ± standard deviation.

Stream	**Unit**	**Paper**	**TechUPGRADE model**	**Note**
Condenser liquid level	%	50.00 ± 0.29	50.00 ± 0.00	
Condenser pressure	kPa	101.28 ± 0.45	101.30 ± 0.27	
Reboiler liquid level	%	49.56 ± 6.12	49.60 ± 0.00	
Feed mass flow	kg/h	3823 ± 511	3883 ± 270	
Top outlet mass flow	kg/h	6866 ± 755	3486 ± 29	vapour to the condenser
Net mass flow, main tower	kg/h	2806 ± 260	1953 ± 291	read as the bottoms
HX, overhead (paper: 'at reboiler')	–	0.953 ± 0.005	0.945 ± 0.047	paper's labels swapped
HX, bottoms (paper: 'at top outlet')	–	0.055 ± 0.004	0.022 ± 0.020	paper's labels swapped
Feed HX mole fraction	–	0.501	0.501 ± 0.003	
Feed tray temperature	°C	76.84 ± 0.33	87.99 ± 2.58	bubble point at 105 kPa
Main tower pressure	kPa	104.43 ± 0.28	105.58 ± 0.27	
Bottom tower pressure	kPa	110.49 ± 0.28	110.44 ± 0.27	
Top tower pressure	kPa	67.31 ± 0.28	101.80 ± 0.27	paper's is below the condenser
Reflux ratio (L/D)	–	0.81 ± 0.16	0.81 ± 0.09	
Reboiler duty	kW	0	370 ± 3	paper: 0 (const.)

Table 2: HX mole-fraction error of the paper's models (its Table 2) and of TechUPGRADE's. The paper predicts its own HYSYS data, TechUPGRADE its own plant model's: different tasks, shown side by side for scale only.

Model	**RMSE**	**R^2**	**Test data**
LSTM	0.00366	0.9482	paper's HYSYS test set
Vanilla MLP	0.00295	0.9645	paper's HYSYS test set
GRU	0.00281	0.9688	paper's HYSYS test set
Transformer	0.00258	0.9721	paper's HYSYS test set
DeepONet	0.00226	0.9796	paper's HYSYS test set
PINN (paper)	0.00143	0.9887	paper's HYSYS test set
TechUPGRADE reduced physics	0.04451	–	paper's transient, held out
TechUPGRADE pHNN (PINN)	0.00798	–	paper's transient, held out

Table 3: TechUPGRADE's PINN on the paper's transient, held out, open loop: the reduced physics it starts from against the trained model.

Quantity	**Reduced physics**	**Trained PINN**
x_D RMSE (mole fraction)	0.0445	0.0080
x_B RMSE (mole fraction)	0.0042	0.0007
x, all stages, RMSE	0.0704	0.0050
T, all stages, RMSE (K)	2.26	0.17
conservation closure (relative)	3.1e-07	2.1e-06

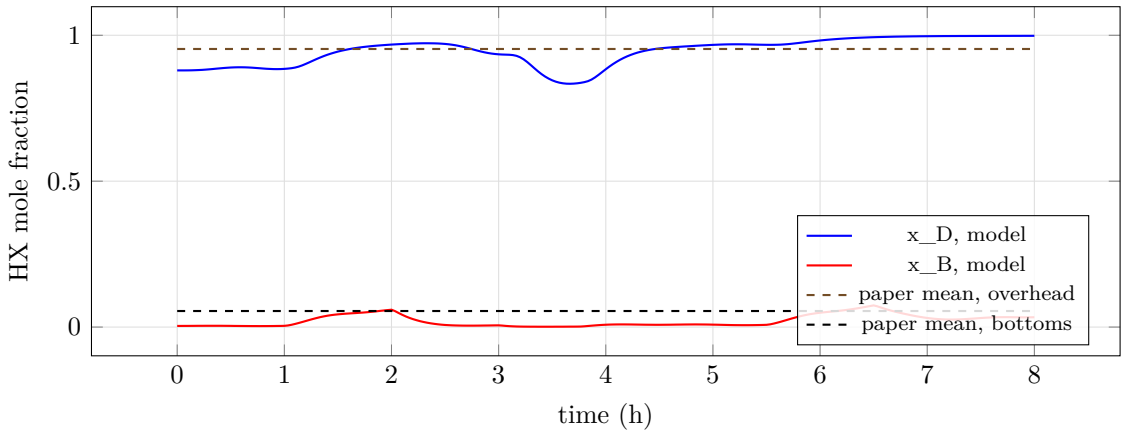


Figure 1: Overhead and bottoms hexane over the paper's 8-hour transient (TechUPGRADE model), with the paper's means (dashed; its labels swapped back).

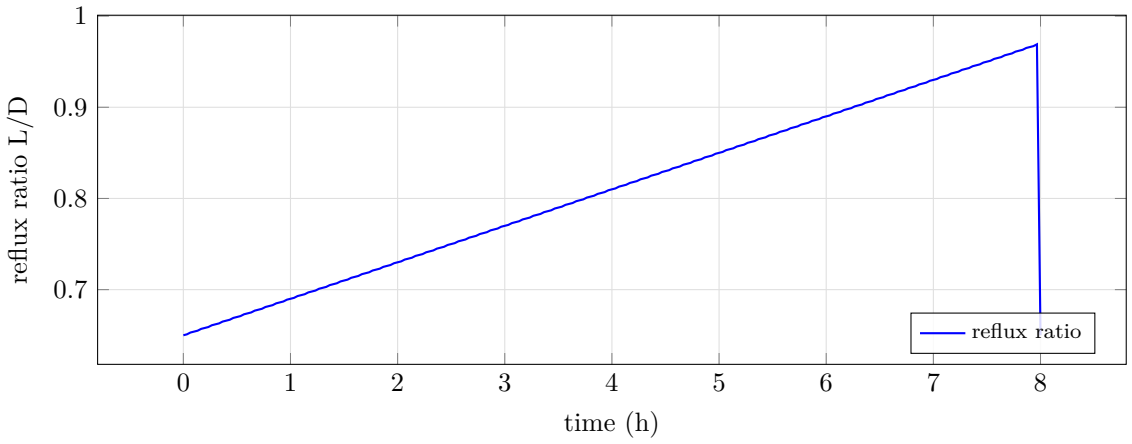


Figure 2: Inputs of the transient: the reflux ratio rising 0.65 to 0.97 and the feed steps.

5 Discussion and Limitations

The reproduction confirms that the PINN framework of Patra et al. (2026) can be implemented in the TechUPGRADE platform, with physics residuals enforcing modified Raoult's law and tray-wise balances, and conservation laws built into the pHNN structure. The trained model achieves lower RMSE than the reduced-physics baseline for all mole-fraction and temperature quantities, demonstrating the value of physics constraints even when trained on platform-generated synthetic data.

However, several limitations arise from differences in data origin and calibration assumptions. The paper's Aspen HYSYS data and the platform's MESH model, while both high-fidelity, differ in numerical implementation, control strategies, and boundary conditions (e.g., reboiler duty, pressure profile, feed tray conditions). The paper's reported top tower pressure (67.31 kPa) is inconsistent with condenser pressure (101.28 kPa), suggesting a possible labeling or reporting error; the model instead follows a physically consistent pressure gradient. Similarly, the paper's feed tray temperature (76.84 °C) is below the bubble point at the reported pressure, whereas the model enforces saturation. The mismatch in overhead mass flow and reboiler duty further indicates differences in how reflux and vapor flow are balanced.

Additionally, the pHNN uses 18 stages (including reboiler and condenser as separate entities), while the paper uses 16 trays; this structural difference may affect convergence and error propagation. The evaluation is performed on the paper's transient, but the model was not trained on that data—only on platform-generated step runs—so the test set is not identical to the paper's. Finally, the paper's reported R^2 for the pHNN is not provided in Table 2, so only RMSE is reported for the TechUPGRADE model.

6 Conclusion

This work reports the reproduction of the PINN digital twin for dynamic tray-wise distillation column modeling on the TechUPGRADE platform. The 16-tray HX/TX column is implemented with identical transient perturbations, and a port-Hamiltonian PINN is trained on platform-generated synthetic data, enforcing the same thermodynamic constraints as in the source study. Evaluation on the paper's held-out transient shows the pHNN achieves an HX mole-fraction RMSE of 0.00798, compared to the paper's reported 0.00143 on its own data. The discrepancy is attributed to

differences in data origin (MESH vs. Aspen HYSYS), model structure (18-stage pHNN vs. 16-tray PINN), and calibration assumptions (e.g., pressure profile, reboiler duty, feed tray conditions). Despite these differences, the model demonstrates improvement over a reduced-physics baseline and retains conservation by construction. The results validate the feasibility of reproducing the PINN framework in TechUPGRADE and highlight the importance of careful calibration and thermodynamic consistency in digital twin development. Future work will explore fine-tuning on Aspen-simulated data and direct comparison of transient responses under identical perturbations.

References

[1] Debadutta Patra, Ayush Bardhan Tripathy, Soumya Ranjan Sahu, Sucheta Panda. *Physics-Informed Neural Network Digital Twin for Dynamic Tray-Wise Modeling of Distillation Columns under Transient Operating Conditions.* arXiv preprint, 2026. `https://arxiv.org/abs/2603.24644`

A Process record

The way from the search that found the source paper to this report, from the platform's records (UTC). The time before a decision at a gate is people's, the time until the implementation (here the reproduction) is on staging is development's, and the rest is the owner's and the platform's (searching, judging, drafting, writing).

Table 4: Process timeline: from the search to the report.

Step	When (UTC)	Took	By	Wait
Search on arXiv: "physics-informed digital twin distillation column tray-wise"	2026-10-24 18:48:05	0 s	techupgrade-admin	machine
The paper found and judged by the owner	2026-10-24 18:49:16	1 min 11 s	owner	machine
The proposal drafted	2026-10-24 18:49:33	17 s	owner	machine
Approved	2026-10-24 18:49:49	16 s	techupgrade-admin	people
Handed to development	2026-10-24 18:49:49	0 s	l2	machine
Implemented / reproduced, on staging	2026-10-24 18:49:56	7 s	l2	development
Release approved	2026-10-24 18:50:40	44 s	techupgrade-admin	people
In production	2026-10-24 18:50:47	7 s	techupgrade-admin	people
The report's first version written	2026-10-24 18:52:33	1 min 47 s	owner	machine
Total		**4 min 29 s**		
of it, people's decisions		1 min 07 s		
of it, development		7 s		
of it, the owner and the platform		3 min 15 s		