Parametric Differentiable Flow Modeling for Fluid Production Optimization Under Flow Assurance Constraints

Luis Kin Miyatake[1,2], Eduardo Camponogara[2], Alexey Pavlov[1]

Abstract— **This work presents a Physics-Informed Neural Network (PINN) framework for multiphase steady-state production systems and its application to constrained production optimization. The PINN is constructed to approximate the coupled momentum and energy equations of a flowline–riser configuration with state-dependent thermophysical properties obtained from tabulated PVT (Pressure-Volume-Temperature) data via differentiable bilinear interpolation. The network is trained by enforcing the governing equations and boundary conditions, yielding continuous pressure and temperature fields along the spatial domain *without requiring labeled data*. Once trained, the neural network weights are frozen and embedded into a production optimization problem, where choke settings and well activation decisions are optimized under global water-handling limits and local operability constraints, including bottom-hole pressure requirements and hydrate avoidance conditions evaluated along the entire flow path. The problem is cast as a mixed-integer nonlinear program, with physical feasibility enforced through a differentiable PINN surrogate.**

I. Introduction

Steady-state multiphase flow modeling in production pipelines plays a central role in the supervision and optimization of oil and gas systems. The governing equations consist of coupled nonlinear conservation laws for mass, momentum, and energy, complemented by closure relations such as equations of state and drift-flux formulations. The interaction among these relations produces complex spatial dynamics that are challenging to resolve efficiently.

Conventional mechanistic simulators [1] typically rely on spatial discretization and stepwise integration procedures. In steady-state scenarios, the problem is often recast as an initial value formulation and solved using shooting strategies to enforce boundary conditions at both ends of the pipeline. Although accurate, such procedures may involve repeated integrations and iterative adjustments, leading to significant computational effort.

Recently, Physics-Informed Neural Networks (PINNs) [2] have emerged as an alternative computational framework that embeds physical laws directly into neural network training. Rather than repeatedly solving the governing equations for each operating condition, the trained network represents a differentiable surrogate of the distributed model. This enables fast evaluations and gradient-based analyses under varying operational scenarios.

From an operational standpoint, the differentiable structure of a trained PINN is particularly attractive for optimization purposes. Instead of being used primarily for parameter inference, the physics-informed surrogate can be embedded directly within gradient-based optimization frameworks. Because the governing equations are enforced during training, the resulting model preserves physical consistency while remaining computationally efficient and fully differentiable with respect to operational variables.

This property enables the formulation of production optimization problems in which control variables—such as choke settings, bottom-hole pressure, or gas injection rates—are adjusted subject to physical and operational constraints. In particular, flow assurance limits, including pressure, temperature, or phase-behavior-related restrictions, can be incorporated explicitly as inequality constraints within the optimization problem. The differentiability of the PINN allows the direct computation of gradients required by nonlinear programming solvers, avoiding repeated numerical integration of the distributed flow model and significantly reducing computational cost.

As a result, PINN-based pipeline models can serve as physics-consistent, differentiable surrogates suitable for production optimization under capacity limitations and flow assurance constraints.

II. Physics-Informed Formulation

Multiphase flow in production systems is characterized by strong nonlinear coupling between phases, governing equations, and state-dependent thermophysical properties. The problem is formulated as a Two-Point Boundary Value Problem (TPBVP), where pressure and temperature evolve along the spatial coordinate under prescribed inlet and outlet conditions.

A. Governing Structure

The distributed model is expressed in one-dimensional form as a coupled system of momentum and energy equations,

$$\begin{cases} \dfrac{\partial P}{\partial x} = f_m(P, T, \mathbf{z}), \\ \dfrac{\partial T}{\partial x} = f_e(P, T, \mathbf{z}), \end{cases} \tag{1}$$

where $f_m(\cdot)$ accounts for hydrostatic and frictional effects, and $f_e(\cdot)$ includes convective heat transfer, gravitational, and Joule–Thomson contributions. The vector $\mathbf{z}$ contains state-dependent variables (e.g., density, viscosity, and solution gas–oil ratio), obtained from thermodynamic closure relations.

[1]Norwegian University of Science and Technology (NTNU), Norway, e-mails: luiskm@stud.ntnu.no, alexey.pavlov@ntnu.no

[2]Federal University of Santa Catarina (UFSC), Brazil, e-mail: eduardo.camponogara@ufsc.br

Thermophysical properties are computed from PVT relations,

$$\mathbf{z} = \phi(P, T), \tag{2}$$

introducing additional nonlinear coupling between pressure and temperature fields.

It is worth emphasizing that the neural network outputs pressure and temperature as explicit functions of the spatial coordinate, thereby capturing their spatial distribution along the flow domain. Since the thermodynamic properties depend on the state variables (P, T), they also vary spatially in a consistent and coupled manner. This represents a fundamental distinction from many classical PINN applications, which treat physical parameters as fixed constants, neglecting their intrinsic dependence on the evolving thermodynamic states.

Closure of the multiphase model is provided by a pattern-free drift-flux [3] formulation, which introduces empirically calibrated algebraic relationships between phase velocities and mixture variables. This closure eliminates the need for explicit flow regime tracking and renders the governing conservation equations mathematically well-posed and fully determined.

B. Boundary Conditions

The spatial domain $x \in [0, L]$ is subject to two-point boundary conditions,

$$\begin{cases} T(0) = T_{\text{in}}, \\ P(L) = P_{\text{out}}, \end{cases} \tag{3}$$

defining a TPBVP (two-point boundary value problem) in which upstream inflow and downstream pressure are prescribed. The boundary conditions are illustrated in Figure 2.

C. Embedding PVT Tables

Thermophysical properties are obtained from thermodynamic equilibrium calculations and are typically generated for industrial applications using dedicated software packages (e.g., Multiflash). These calculations provide tabulated PVT data defined on a structured grid in the (P, T) space. In this work, such tables are treated as the thermodynamic input to the multiphase flow model.

Let (P_i, P_{i+1}) and (T_j, T_{j+1}) denote the neighboring grid points such that

$$P_i \le P \le P_{i+1}, \qquad T_j \le T \le T_{j+1}.$$

The four surrounding tabulated values, as shown in Figure 1, are

$$\phi_{i,j}, \quad \phi_{i+1,j}, \quad \phi_{i,j+1}, \quad \phi_{i+1,j+1}.$$

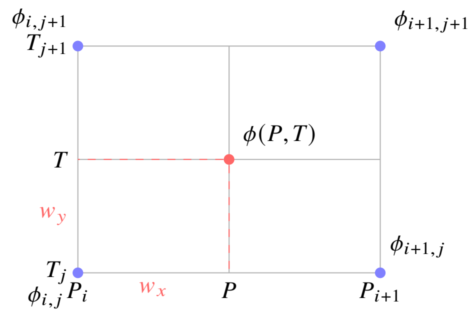

Fig. 1. Schematic of the bilinear interpolation used to evaluate the property $\phi(P, T)$ from the tabulated PVT data.

Normalized interpolation weights are defined as

$$w_x = \frac{P - P_i}{P_{i+1} - P_i}, \qquad w_y = \frac{T - T_j}{T_{j+1} - T_j}. \tag{4}$$

The property value at (P, T) is then obtained via bilinear interpolation,

$$\begin{aligned} \phi(P,T) = \; & (1 - w_x)(1 - w_y)\, \phi_{i,j} + w_x(1 - w_y)\, \phi_{i+1,j} \\ & + (1 - w_x) w_y\, \phi_{i,j+1} + w_x w_y\, \phi_{i+1,j+1}. \end{aligned} \tag{5}$$

Since w_x and w_y are continuous functions of P and T inside each interpolation cell, the interpolated property $\phi(P, T)$ remains locally differentiable with respect to the neural-network outputs and weights. Therefore, automatic differentiation can propagate gradients through the PVT interpolation step during PINN training.

At cell boundaries, however, bilinear interpolation is generally not differentiable, since the local interpolation derivatives may change discontinuously from one cell to the next. In practice, automatic differentiation uses the local derivative associated with the active interpolation cell. Since the PVT tables are defined on a fine pressure–temperature grid, these derivative jumps are expected to be small, reducing the effect of transitions between cells during Adam and L-BFGS training.

D. Physics-Informed Neural Networks Modeling

Instead of solving the TPBVP using marching or shooting methods, a physics-informed neural network $\mathcal{N}_\phi$ is introduced to approximate the continuous mappings,

$$(P, T) = \mathcal{N}_\phi(x, \theta, u), \tag{6}$$

where θ represents operational parameters (e.g., BSW (water cut)[1], and GOR), and u are the manipulated variables (e.g., choke settings).

The network is trained by minimizing a scaled composite loss function,

$$\mathcal{L} = \mathcal{L}_{\text{momentum}} + \mathcal{L}_{\text{energy}} + \mathcal{L}_P + \mathcal{L}_T, \tag{7}$$

[1] Water cut is the fraction of water in the total liquid production stream, typically expressed as a percentage.

where the governing equation residuals are defined as

$$\mathcal{L}_{\text{momentum}} = \left\| \frac{1}{\tau_{\text{mom}}} \left(\frac{\partial P}{\partial x} - f_m(P, T, \mathbf{z}) \right) \right\|^2, \tag{8}$$

$$\mathcal{L}_{\text{energy}} = \left\| \frac{1}{\tau_{\text{energy}}} \left(\frac{\partial T}{\partial x} - f_e(P, T, \mathbf{z}) \right) \right\|^2. \tag{9}$$

The boundary condition losses are given by

$$\mathcal{L}_P = \left\| \frac{P(x = L) - P_{\text{out}}}{\tau_P} \right\|^2, \tag{10}$$

$$\mathcal{L}_T = \left\| \frac{T(x = 0) - T_{\text{in}}}{\tau_T} \right\|^2 \tag{11}$$

where τ_{mom}, τ_{energy}, τ_P, and τ_T are scaling parameters that normalize the magnitude of each residual term, mitigating imbalance due to heterogeneous physical scales and improving training stability.

In conventional neural network training, a data loss term is typically included to enforce agreement with measured observations. However, in this work we adopt a parametric PINN strategy (Eq. 6), in which the parameters θ are defined over the entire operational range of the well.

As a result, the network is trained to represent a family of steady-state solutions parameterized by θ, enabling its use across multiple operating scenarios without retraining. This parametric formulation is particularly suited for optimization purposes, since the production optimization problem introduced in Section III seeks the optimal control variables $\mathbf{u}$ over the full admissible operational parameters θ captured by the parametric PINN.

III. Production Optimization Formulation

A. Problem Overview

We consider the optimal production allocation problem [4], [5] of a multi-well system operating under global water-handling limitations and local operability constraints. Each well is connected to a shared production system and must satisfy upstream pressure constraints and hydrate-avoidance conditions along the flowline.

The optimization problem simultaneously determines:

- Continuous choke variables (u),
- Binary variables representing well open/shut decisions,
- Nonlinear well responses predicted by a physics-informed neural network.

Due to the presence of discrete decisions and nonlinear physical models, the resulting problem is formulated as a Mixed-Integer Nonlinear Program (MINLP).

B. Decision Variables and Differentiable Production Model

For each well $i = 1, \ldots, N$, we define:

$$\begin{cases} u_i \in [\Delta P_{\min}, \Delta P_{\max}] & \text{choke pressure drop,} \\ t_i \in \{0, 1\} & \text{well status indicator.} \end{cases} \tag{12}$$

The choke pressure drop u_i represents the pressure drop across the choke valve. The binary variable t_i allows the optimizer to deactivate wells when required by feasibility constraints.

Well behavior is predicted by a differentiable surrogate model represented by a physics-informed neural network, which approximates the continuous spatial fields along the production line.

The network defines the mappings

$$(P_i(x), T_i(x)) = \mathcal{N}(x, \text{BSW}_i, u_i), \quad x \in [0, L]. \tag{13}$$

Thus, for any spatial coordinate x, the PINN returns the corresponding pressure and temperature values.

It is important to clarify that the primary focus of this work is the *production optimization problem*. In this setting, the decision variables are the control inputs u_i, while the parameters BSW_i are treated as fixed and known quantities. The upstream pressure is obtained as the inlet evaluation of the field,

$$P_{\text{inlet},i} = P_i(x = 0, \text{BSW}_i, u_i). \tag{14}$$

For constraint enforcement along the flowline, the continuous fields are evaluated at a finite set of spatially distributed points

$$x_k, \quad k = 1, \ldots, K,$$

representing discrete locations along the wellbore or pipeline.

Accordingly, we define

$$P_{i,k} = P_i(x_k), \qquad T_{i,k} = T_i(x_k), \tag{15}$$

which are used to enforce hydrate and operability constraints at multiple spatial positions.

The total liquid rate of well i is computed via a linear Inflow Performance Relation (IPR),

$$Q_i = \text{PI}_i \left(P_{\text{static},i} - P_{\text{inlet},i} - u_i \right). \tag{16}$$

Here, $P_{\text{static},i}$ denotes the static pressure of well i. The parameter PI_i is the productivity index.

Oil and water production rates are defined as

$$\begin{cases} Q_{o,i} = t_i(1 - \text{BSW}_i)Q_i, \\ Q_{w,i} = t_i \text{BSW}_i Q_i. \end{cases} \tag{17}$$

The multiplicative structure ensures that production is automatically set to zero when a well is shut ($t_i = 0$).

C. Objective Function

The objective is to maximize total oil production while allowing wells to be selectively shut if local constraints cannot be satisfied.

The optimization problem is therefore formulated as

$$\max_{\boldsymbol{u}, \boldsymbol{t}} \sum_{i=1}^{N} Q_{o,i}. \tag{18}$$

D. Operational Constraints

1) Global Water Handling Constraint: The total produced water must not exceed the facility handling capacity:

$$\sum_{i=1}^{N} t_i \, \text{BSW}_i \, Q_i \leq Q_w^{\max}. \tag{19}$$

2) Upstream Pressure Constraint: Each producing well must satisfy a minimum pressure:

$$P_{\text{inlet},i} + u_i \geq P_i^{\min} - M_{\text{P}}(1 - t_i), \qquad \forall i. \tag{20}$$

where M_{P} is a sufficiently large constant (big-M).

3) Hydrate Operability Constraint: Gas hydrates are crystalline solid compounds formed when water and light gas molecules, such as methane, coexist under high-pressure and low-temperature conditions. In offshore production systems, hydrate formation is a major flow-assurance concern because solid deposits may restrict or block the flowline.

Therefore, hydrate operability constraints are commonly included to ensure that the predicted pressure–temperature trajectory remains outside, or sufficiently close to outside, the hydrate stability region (see, e.g., [6], [7]). The hydrate stability region (that should be avoided) corresponds to the region located to the left of the dashed black hydrate-equilibrium curve shown in Figure 5.

To prevent hydrate formation, hydrate operability is enforced along the flowline using a logarithmic equilibrium relation:

$$T_{i,k} \geq \left(A + B\log_{10} P_{i,k}\right) - \text{tol}_T - M_{\text{hyd}}(1 - t_i), \forall i, k, \tag{21}$$

where M_{hyd} is a sufficiently large constant (big-M).

Here, the index $k = 1, \ldots, K$ denotes discrete spatial locations x_k along the wellbore or flowline, where hydrate risk is evaluated. The quantities $P_{i,k} = P_i(x_k)$ and $T_{i,k} = T_i(x_k)$ are obtained from the continuous PINN-predicted fields.

The term $A + B\log P_{i,k}$ defines the logarithmic hydrate equilibrium curve (dashed black line in Figure 5), which determines the temperature below which hydrates may form at pressure $P_{i,k}$. The coefficients A and B are obtained by fitting the predictions of a thermodynamic model for the system of interest. The resulting expression provides a compact representation of the hydrate stability boundary for use in the optimization framework.

Operating below this boundary implies entering the hydrate stability region. The tolerance parameter tol_T controls how far the computed temperature profile is allowed to fall below this curve. Larger values of tol_T permit operation deeper inside the hydrate region, increasing the risk of flow blockage.

E. Final MINLP Formulation

The complete production optimization problem is cast as:

$$\begin{aligned}
\max_{\boldsymbol{u},\boldsymbol{t}} \quad & \sum_{i=1}^{N} t_i(1 - \text{BSW}_i)Q_i \\
\text{s.t.} \quad & \sum_{i=1}^{N} t_i \text{BSW}_i Q_i \leq Q_w^{\max} \\
& P_{inlet,i} \geq P_i^{\min} - u_i - M_{\text{P}}(1 - t_i), \quad \forall i \\
& T_{i,k} \geq (A + B\log_{10} P_{i,k}) - \text{tol}_T - M_{\text{hyd}}(1 - t_i), \forall i, k \\
& u_i \in [\Delta P_{\min}, \Delta P_{\max}], \quad t_i \in \{0, 1\}, \quad \forall i.
\end{aligned} \tag{22}$$

The differentiable PINN surrogate provides continuous spatial fields $P_i(x)$ and $T_i(x)$ along the entire production system. Rather than imposing constraints at a single operating point, the optimization framework leverages these predicted pressure and temperature profiles to enforce operability conditions distributed along the full flow path.

In particular, hydrate and pressure constraints are evaluated at multiple spatial locations, ensuring compliance throughout the pipeline.

This formulation exploits the spatially resolved structure of the PINN, transforming distributed physical behavior into optimization-ready constraints. This allows spatially distributed physical constraints and integer well-status decisions to be treated consistently within a single optimization problem.

F. Numerical Optimization Framework

The mixed-integer nonlinear programming problem is solved using BONMIN [8], which combines Branch-and-Bound (B&B) strategies with Nonlinear Programming (NLP) subproblem solvers. In particular, continuous relaxations arising within the B&B search are handled using IPOPT [9], an interior-point method for large-scale nonlinear optimization.

The optimization problem is implemented in CasADi [10], which provides symbolic representation and Automatic Differentiation (AD) capabilities. This enables exact gradient and Jacobian evaluation for both the objective function and the constraints, significantly improving numerical robustness and convergence performance.

It should be noted that BONMIN is a practical solver choice for the present formulation because the problem contains both binary well-status variables and nonlinear differentiable constraints provided by the PINN surrogate. However, since the resulting MINLP is nonconvex, global optimality cannot in general be guaranteed by BONMIN/IPOPT. Therefore, the obtained solution should be interpreted as the best solution found by the selected solver under the specified initialization, bounds, and tolerances.

Deterministic global MINLP solvers, such as Couenne [11] or SCIP [12], could in principle be considered to provide stronger global optimality guarantees. However, their use would typically require an explicit algebraic transcription of the neural-network surrogate into the optimization problem, introducing auxiliary variables and equality constraints associated with the neurons and layers of the network, thereby substantially increasing the size and complexity of the MINLP.

In contrast, the trained PINN is embedded in CasADi as an explicit differentiable mapping. CasADi evaluates the objective and nonlinear constraints through the network and uses automatic differentiation to provide the gradients and Jacobians required by IPOPT within the BONMIN framework. As a result, the number of decision variables scales linearly with the number of wells ($O(N)$), remains independent of the neural-network architecture, and avoids the much larger algebraic reformulation required by deterministic global MINLP solvers.

In practice, solution robustness was further assessed by solving the optimization problem from multiple initial guesses for the continuous decision variables. In the present case study, the BONMIN-based strategy consistently converged to the same operating point across the tested initializations. This observation increases confidence in the obtained solution, although it does not constitute a formal proof of global optimality.

G. From Simulation-Based Workflows to PINN-Based Optimization

Traditional production-optimization workflows commonly rely on mechanistic simulators to generate tabulated data over a predefined range of operating conditions. These tables are then used to construct sensitivity analyses, lookup models, or surrogate representations for optimization. However, the tabulated data must be regenerated whenever the model parameters are updated, making the computational cost of maintaining an optimization-ready model high.

This repeated regeneration of tabulated data is often the main computational bottleneck, since it requires solving the governing equations many times over the relevant operational domain. A trained PINN directly addresses this bottleneck because new operating conditions can be evaluated through a single neural-network forward pass. In our previous work [13], this type of evaluation was shown to be several hundred times faster than conventional numerical solvers.

This observation enables two complementary uses of the PINN framework. First, the trained PINN can be used as a fast data generator to replace the mechanistic simulator in traditional lookup-table-based workflows. In this case, the same optimization methodologies currently employed can be preserved, while the computational cost associated with generating tabulated data is reduced substantially.

Second, once the network is trained, its weights are frozen, and the PINN itself becomes a *physics-consistent differentiable surrogate model* over the admissible ranges of the operational parameters θ and manipulated variables u. The trained network can then be embedded directly into the optimization problem, where automatic differentiation provides exact derivatives of the objective and constraints with respect to the optimization variables.

This second use represents the strategy adopted in the present work. The PINN is incorporated explicitly into a mixed-integer nonlinear programming (MINLP) formulation, allowing distributed pressure and temperature constraints to be enforced along the spatial domain *without requiring additional surrogate construction or numerical integration during optimization*.

Therefore, the proposed framework offers both a computationally efficient replacement for traditional simulation-based data generation and a fully differentiable surrogate model that can be integrated directly into mathematical programming algorithms.

IV. Problem Description: Subsea Flow Pipelines

Figure 2 illustrates the production system considered in this study. The configuration consists of a 3000 m horizontal flowline followed by a 2000 m vertical riser. The horizontal segment is exposed to the surrounding seabed temperature of 5 °C, which induces progressive cooling of the flowing mixture as heat is exchanged with the external environment.

The inlet temperature is prescribed as $T(x = 0) = T_{\text{in}}$, while the outlet pressure is fixed at $P(x = L) = P_{\text{out}}$. These conditions define a two-point boundary value problem along the spatial coordinate $x \in [0, L]$.

The mass source of the system is governed by an inflow performance relation, as defined in Equation 16.

The manipulated variable in this problem is the choke pressure drop, denoted as $u = \Delta P_{\text{choke}}$, which directly controls the production rate.

Along the spatial domain, pressure $P(x)$ and temperature $T(x)$ are computed continuously, and state-dependent thermodynamic properties $\phi(P, T)$ (e.g., density, viscosity, solubility ratio) are evaluated as functions of the local state. This configuration represents a strongly coupled nonlinear system in which model parameters vary with the thermodynamic state of the fluid.

The dependence of thermophysical properties on (P, T) is obtained through bilinear interpolation of tabulated PVT data, as discussed in Section II-C. Therefore, state-dependent parameters are imposed along the entire pipeline.

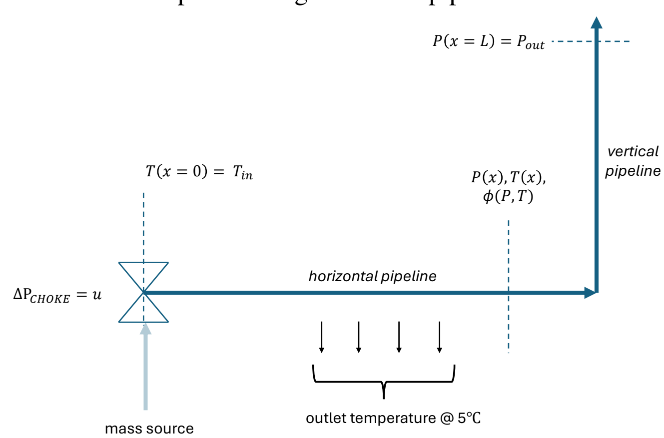

Fig. 2. Flow system with horizontal flowline and vertical riser, subject to inlet temperature and outlet pressure boundary conditions and choke-controlled production.

V. Results

A. Validation of the PINN model with Traditional Solvers

The Figure 3 compares the classical marching method (solid lines) and the PINN-based solution (dashed lines) for different water cut (BSW) values.

The pressure–temperature profiles predicted by the PINN closely match those obtained from the traditional marching method. The agreement demonstrates that the PINN model preserves the underlying physical behavior of the multiphase system while accurately reproducing the steady-state solution across a family of parameter (BSW) values.

Table I shows that the mean absolute percentage error (MAPE) between the traditional marching solution and the

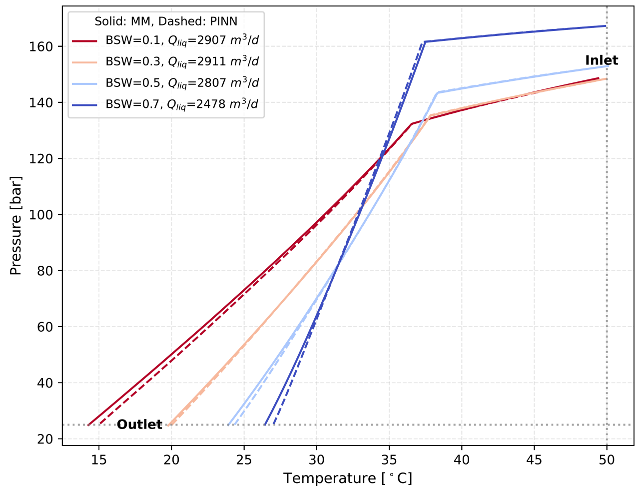

Fig. 3. Comparison between the classical marching method (solid lines) and the PINN-based solution (dashed lines) for different water cut (BSW) values. Each BSW leads to a distinct equilibrium liquid rate, with higher BSW resulting in lower production rate due to increased hydrostatic effects. The PINN accurately reproduces the pressure–temperature profiles while satisfying the prescribed boundary conditions ($T_{\text{in}} = 50°\text{C}$, $P_{\text{out}} = 25$ bar). The change of slope in the pressure profile reflect the transition from the horizontal flowline to the vertical riser section.

PINN predictions, which are generally below 1.0%. MAPE is defined as

$$\text{MAPE} = \frac{1}{N_x} \sum_{j=1}^{N_x} \left| \frac{y_{\text{PINN}}(x_j) - y_{\text{ref}}(x_j)}{y_{\text{ref}}(x_j)} \right|, \tag{23}$$

where N_x denotes the number of spatial evaluation points along the profile, y_{PINN} represents the PINN prediction, and y_{ref} corresponds to the reference solution obtained using the marching method.

TABLE I

Mean Absolute Percentage Error (MAPE %) for PINN pressure and temperature predictions.

BSW	P [%]	T [%]
0.1	0.106	0.677
0.3	0.044	0.456
0.5	0.139	0.463
0.7	0.087	0.560
0.9	0.279	1.393

B. Optimization Problem Setup

The production optimization problem is formulated considering operational and flow assurance constraints representative of realistic offshore conditions. A maximum liquid handling capacity of 13,000 m^3/d is imposed at the facility level. The water cut (BSW) is treated as a parametric variable ranging from 0.1 to 0.95, with discrete increments of 0.1, defining the operational envelope explored in the optimization study.

Hydrate formation risk is incorporated through an empirical hydrate equilibrium curve parameterized by constants A and B, defining a temperature–pressure relationship that constrains the admissible operating region (Equation 21). The hydrate constraint is enforced along the entire spatial profile of the pipeline, ensuring thermodynamic feasibility under multiphase conditions.

A fundamental trade-off arises between production maximization and hydrate risk tolerance, tol_T. If the hydrate constraint is relaxed (i.e., larger safety margin), the optimization problem becomes effectively unconstrained with respect to flow assurance, allowing the objective function, defined here as oil production rate, to reach its maximum value limited only by capacity constraints. Conversely, if the hydrate tolerance is tightened, the constraint becomes active over larger portions of the spatial domain and for a greater number of wells, restricting admissible operating conditions and reducing the achievable oil production rate.

This interaction generates a Pareto frontier (Figure 4) that characterizes the balance between maximizing oil production and limiting hydrate formation risk.

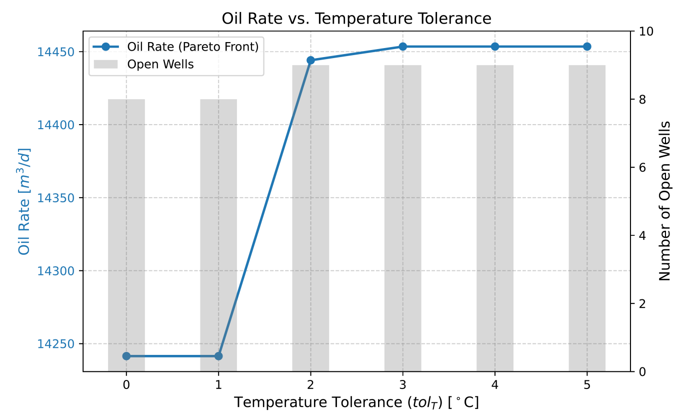

Fig. 4. Total oil production rate as a function of hydrate temperature tolerance, illustrating the Pareto trade-off between production maximization and flow assurance risk. An optimal compromise is observed at a tolerance of 2°C (Pareto elbow), where oil production increases significantly while maintaining acceptable hydrate safety margins. For lower tolerances (0–1°C), the hydrate constraint becomes active over larger portions of the spatial domain, forcing the optimizer to operate with only eight wells open in order to satisfy the imposed flow assurance constraints (Equation 21), thereby reducing total oil production. Although further relaxation of the tolerance beyond 2°C yields marginal production gains, the incremental benefit does not justify the increased operational risk associated with operating closer to hydrate formation conditions, which may lead to severe flow assurance issues.

C. Operating Point Analysis at the Pareto Elbow

Figure 5 presents the pressure–temperature operating profiles for the nine wells at the Pareto elbow ($tol_{\text{T}} = 2°\text{C}$), together with the hydrate equilibrium curve. Wells are indexed in increasing order of water cut (BSW), such that lower-indexed wells correspond to lower BSW values. These wells exhibit higher production rates and consequently operate farther from the hydrate envelope due to reduced hydrostatic effects and more favorable thermodynamic conditions.

As the BSW increases, the operating profiles move closer to the hydrate equilibrium boundary. In particular, Well 9 operates effectively at the hydrate constraint limit. The active constraint occurs near the base of the riser section, where the pressure is highest and the temperature approaches the prescribed hydrate tolerance (tol_T) of 2°C. This location corresponds to the most critical point along the spatial domain in terms of hydrate formation risk.

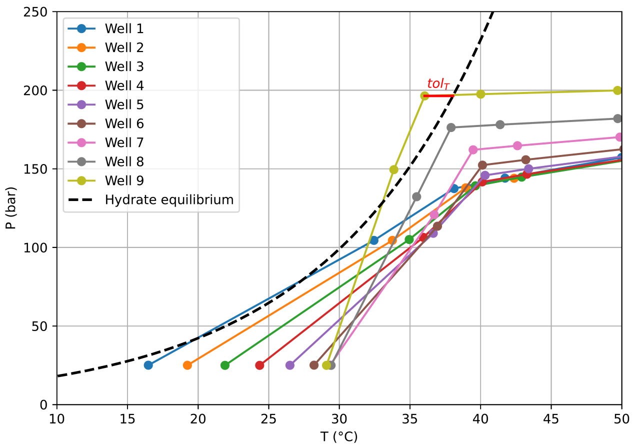

Fig. 5. Pressure–temperature operating profiles of the nine wells at the Pareto elbow ($tol_T = 2$°C) together with the hydrate equilibrium curve (dashed line). Operating conditions located to the right of the equilibrium curve (higher temperatures and lower pressures) correspond to hydrate-safe regions, whereas points approaching or crossing the curve indicate hydrate formation risk. Lower-indexed wells (lower BSW) operate farther from the hydrate envelope due to higher production rates, while higher-BSW wells move closer to the constraint boundary. Well 9 reaches the active hydrate constraint near the base of the riser, where the minimum safety margin is attained. This critical location is highlighted in red.

VI. Conclusions and Future Work

This work presented a structured integration between physics-informed neural networks and constrained production optimization. The main methodological contributions can be summarized as follows:

- *Parametric PINN formulation for operational optimization:* A parametric steady-state PINN was trained over the admissible operational domain, enabling the modeling of a family of distributed solutions. This formulation allows operational decision variables (e.g., choke settings) to be directly embedded into the optimization problem without retraining the surrogate model.
- *Differentiable embedding within a MINLP framework:* The trained PINN is incorporated as an explicit differentiable mapping, rather than being expressed as a set of equality constraints. Thus, the number of decision variables scales linearly with the number of wells ($O(N)$), remaining independent of the neural network architecture or spatial discretization. Automatic differentiation through CasADi enables consistent gradient evaluation for both objective and constraint functions.
- *Natural treatment of spatially distributed constraints:* The continuous fields predicted by the PINN allow operability conditions—such as hydrate avoidance—to be enforced along the spatial domain. This enables distributed physical constraints to be integrated seamlessly within a single optimization problem, preserving physical consistency while maintaining computational tractability.

Future work will extend the proposed framework to two-phase flow modeling in networked production systems. In such settings, multiple PINN models must be trained and coupled to represent interconnected pipelines and manifold structures. This modular architecture would enable large-scale optimization problems involving multiple wells and shared infrastructure to be solved efficiently using mathematical programming algorithms. By exploiting the computational efficiency and differentiability of physics-informed neural networks, the framework has the potential to support real-time decision-making in complex field-wide production networks.

Acknowledgment

This research was funded in part by Petróleo Brasileiro S.A. (Petrobras), the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) under grants 308624/2021-1 and 402099/2023-0, and the BRU21 research and innovation program at NTNU (Digital and Automation Solutions for the Offshore Energy Industry).

References

[1] Ovadia Shoham. *Mechanistic Modeling of Gas-Liquid Two-Phase Flow in Pipes*. Society of Petroleum Engineers, Richardson, TX, 2006.

[2] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

[3] S. Bhagwat and A. Ghajar. A flow pattern independent drift flux model based void fraction correlation for a wide range of gas–liquid two phase flow. *International Journal of Multiphase Flow*, 59:186–205, 2014.

[4] E. R. Müller, E. Camponogara, L. O. Seman, et al. Short-term steady-state production optimization of offshore oil platforms: wells with dual completion (gas-lift and ESP) and flow assurance. *TOP*, 30:152–180, 2022.

[5] E. R. Müller, E. Camponogara, L. O. Seman, et al. Optimizing operations sequencing and demulsifier injection in offshore oil production platforms. *IEEE Access*, 11:33414–33427, 2023.

[6] E. Dendy Sloan and Carolyn A. Koh. *Clathrate Hydrates of Natural Gases*. CRC Press, Boca Raton, 3 edition, 2007.

[7] John J. Carroll. *Natural Gas Hydrates: A Guide for Engineers*. Gulf Professional Publishing, 3 edition, 2014.

[8] P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008.

[9] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.

[10] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl. CasADi: A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11:1–36, 2019.

[11] Pietro Belotti, Jon Lee, Leo Liberti, François Margot, and Andreas Wächter. Branching and Bounds Tightening Techniques for Non-Convex MINLP. *Optimization Methods and Software*, 24(4-5):597–634, 2009.

[12] Tobias Achterberg. Scip: Solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.

[13] Luis Kin Miyatake, Eduardo Camponogara, Eric Aislan Antonelo, and Alexey Pavlov. Physics-informed neural networks for control of single-phase flow systems governed by partial differential equations. *Engineering Applications of Artificial Intelligence*, 160:111869, 2025.