

Article

Physics-Informed Neural Network Surrogate Modeling of Three-Dimensional Gas–Liquid Hydrodynamics in Aerated Double-Impeller Stirred Tanks

Forough Sharifi, Ali Lohi  and Farhad Ein-Mozaffari * 

Department of Chemical Engineering, Toronto Metropolitan University, 350 Victoria Street, Toronto, ON M5B 2K3, Canada; forough.sharifi@torontomu.ca (F.S.); alohi@torontomu.ca (A.L.)

* Correspondence: fmozaffa@torontomu.ca

Abstract

This study develops a physics-informed surrogate for predicting three-dimensional gas–liquid hydrodynamics in an aerated double-impeller stirred tank. Reference fields were generated using steady-state computational fluid dynamics (CFD) simulations in ANSYS Fluent at impeller speeds of 200–500 rpm. Three models with identical network architectures were examined: a supervised neural network (NN) trained on CFD data, a boundary-informed neural network (BINN) incorporating boundary constraints, and a physics-informed neural network (PINN) incorporating boundary and reduced physics constraints. At 400 and 500 rpm, the PINN outperformed the NN and BINN, achieving normalized liquid- and gas velocity errors below 10% and pressure errors of approximately 3–4%. The NN and BINN showed greater attenuation of impeller-induced gradients in high-shear regions. At the unseen intermediate speeds of 350 and 450 rpm, the PINN reproduced the principal CFD flow structures. Normalized gas velocity errors were approximately 11–14%, while normalized pressure errors ranged from 4.48% to 9.22%, depending on the operating speed and error metric. Each CFD simulation required approximately 100–140 h, whereas PINN inference at 200,000 spatial locations required approximately 2–3 h per operating condition. The initial cost of generating the CFD training data must be considered when assessing the benefit of repeated surrogate evaluations.

Keywords: physics-informed neural networks; multiphase flow; stirred tanks; surrogate modeling; computational fluid dynamics

1. Introduction

Stirred-tank bioreactors are widely used in chemical and biotechnological industries for mixing, chemical reactions, and fermentation processes [1–4]. They commonly consist of a cylindrical vessel equipped with rotating impellers and stationary baffles to improve mixing and suppress vortex formation. For high-aspect-ratio vessels, multi-impeller configurations are often employed to enhance axial circulation, gas dispersion, and mixing uniformity throughout the tank [5–7]. In aerated bioreactors, uniform gas distribution is particularly important because poor dispersion can reduce oxygen transfer efficiency and create hydrodynamic heterogeneities and dead zones [8]. Double-impeller mixing systems therefore play a critical role in large-scale bioprocessing by improving gas–liquid mass transfer and promoting more homogeneous hydrodynamic conditions along the reactor height [9,10].



Academic Editor: Nuno M.C. Oliveira

Received: 11 August 2026

Revised: 25 September 2026

Accepted: 2 October 2026

Published: 9 October 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

The versatility of stirred-tank bioreactors enables their widespread application in pharmaceutical, biochemical, and fermentation processes; however, this flexibility also introduces significant challenges in reactor design, scale-up, and optimization because operating conditions and geometry strongly influence hydrodynamics and gas–liquid mass transfer [8,11]. In aerated systems, local velocity fields, gas holdup, and mixing intensity directly affect momentum transport and oxygen transfer rates, thereby influencing the local hydrodynamic and mass-transfer conditions experienced by suspended cells and microorganisms [12]. Consequently, there is increasing demand for predictive modeling frameworks capable of rapidly evaluating hydrodynamics for process optimization, scale-up, and process-control applications [13,14].

Computational fluid dynamics (CFD) has become a powerful tool for analyzing multi-phase hydrodynamics in stirred tanks because it provides detailed predictions of velocity, pressure, and gas distribution fields [15,16]. However, high computational cost, mesh sensitivity, turbulence-model dependence, and long simulation times limit the application of CFD for rapid parametric studies, optimization, and real-time process monitoring [13]. These limitations are particularly restrictive when multiple operating conditions or geometrical configurations must be evaluated. To address these challenges, reduced-order and surrogate modeling approaches have been explored as computationally efficient alternatives [17,18].

Purely data-driven surrogate models, including neural network-based approaches, have shown promise for predicting stirred-tank hydrodynamics [10,19,20]. Nevertheless, such models generally require large labeled datasets and often lack physical interpretability and robustness outside the training domain [21–23]. Because these models rely solely on data, they may fail to preserve governing physical behavior when extrapolated to unseen operating conditions.

By including equation residuals and boundary conditions in the training loss, PINNs can guide predictions toward the behavior represented by those constraints. Their accuracy and generalization nevertheless depend on the suitability of the equations, the training data, and the operating conditions evaluated [24–26]. Operator-learning approaches, including DeepONets and Fourier Neural Operators (FNOs), have further improved the representation of PDE-governed systems [27,28]. However, these approaches still depend heavily on large, high-quality datasets and do not explicitly enforce governing physical laws in the loss function, limiting their ability to guarantee physically consistent predictions outside the training regime [29]. To overcome these limitations, physics-informed machine learning frameworks have been developed by incorporating governing equations, conservation laws, and boundary constraints directly into the training process [30,31]. Physics-informed neural networks (PINNs), introduced by Raissi et al. [31,32], incorporate governing PDE residuals, boundary conditions, and physical constraints directly into the neural network loss function, thereby constraining predictions to physically admissible solutions. This formulation improves data efficiency, generalization capability, and robustness while preserving physical consistency. PINNs have been successfully applied to a range of flow and transport problems, including fluid-flow prediction, turbulence modeling, and bioprocess-relevant systems [33–35].

Recent applications of physics-informed machine learning in chemical engineering extend beyond canonical fluid-flow problems. Velioglu et al. [36] applied PINNs to dynamic chemical-process models with incomplete mechanistic knowledge and limited data, including a Van de Vusse continuously stirred-tank reactor and a liquid–liquid separator. In stirred systems, Trávníková et al. [37] developed a hierarchy of PINN formulations for two-dimensional stirred-tank flow and subsequently quantified labeled-data requirements relative to supervised and boundary-informed neural networks. Horiguchi et al. further

demonstrated CFD-informed PINNs for high-resolution prediction of flow-induced stresses relevant to suspension bioreactors. More recently, adaptive weighting and collocation strategies have been proposed specifically for chemical-process PINNs to address imbalance among competing loss terms. Collectively, these studies demonstrate the progression from purely data-driven CFD surrogates toward hybrid models incorporating mechanistic constraints; however, the application of such approaches to fully three-dimensional aerated multi-impeller gas–liquid flow remains comparatively less developed [36].

Despite these advances, existing PINN studies for stirred systems have largely focused on single-phase flows, simplified geometries, and two-dimensional domains [29,37]. In particular, previous work has established comparisons among purely data-driven NNs, boundary-informed models, and PINNs and has examined their data requirements in simplified stirred-tank configurations. Extending such approaches to three-dimensional gas–liquid stirred tanks introduces additional challenges associated with phase coupling, rotating internals, spatially heterogeneous gas distributions, and the computational complexity of incorporating multiphase physics into the learning framework. Consequently, the application of physics-informed surrogate modeling to three-dimensional gas–liquid stirred-tank hydrodynamics remains comparatively limited.

Motivated by this gap, the present study develops and evaluates a physics-informed surrogate framework for predicting three-dimensional hydrodynamics in an aerated double-impeller stirred tank. The methodological contribution lies not in the NN/BINN/PINN comparison itself, but in extending this framework to a three-dimensional gas–liquid system and incorporating a reduced multiphase physics formulation into surrogate training. Three progressively constrained strategies are evaluated: (i) a purely data-driven NN, (ii) a BINN incorporating boundary information, and (iii) a PINN incorporating both boundary conditions and reduced physics-based constraints. Their performance is systematically compared to determine how increasing levels of physical information affect predictive accuracy, robustness, data efficiency, and interpolation at unseen intermediate operating conditions. The influence of labeled-data availability is also examined to assess the practical trade-off between CFD-data requirements and physics-based regularization.

It should be emphasized that the physics incorporated into the PINN constitute a reduced representation of the multiphase system rather than the complete Euler–Euler formulation used in the underlying CFD simulations. The purpose of this reduced formulation is therefore not to reproduce all resolved multiphase mechanisms within the neural network, but to introduce physically relevant constraints that guide surrogate learning while maintaining computational tractability. The following sections describe this formulation, evaluate its predictive performance, and discuss its capabilities and limitations for three-dimensional gas–liquid stirred-tank applications.

Within its present domain, the surrogate provides a rapid mapping from spatial location and impeller rotational speed to the principal steady hydrodynamic fields for a fixed reactor geometry, gas flow rate and fluid system. This capability is relevant to repeated operating-condition screening and can serve as a reduced-order component in future optimization or digital-twin architectures where low-latency flow-field estimates are required. The present model should not, however, be interpreted as a complete digital-twin or a scale-up model because online state updating, uncertainty quantification, closed-loop control, and geometric scaling are not included. Extension to those applications requires additional operating and geometric parameters to be incorporated into the surrogate input space.

2. Surrogate Modeling Framework and CFD Methodology

Prior studies have demonstrated that PINNs outperform purely data-driven neural networks, particularly when labeled data are limited or expensive to obtain [31,38]. Building on these findings, this work develops and evaluates a PINN-based surrogate modeling framework for a three-dimensional gas–liquid double-impeller stirred bioreactor. To assess the benefit of incorporating different levels of physical information, two baseline models were also developed. The first was a purely supervised neural network trained exclusively on labeled CFD data. The second was a BINN, in which boundary conditions were incorporated into the loss function without enforcing the governing equations [37]. The BINN therefore represents an intermediate approach between supervised and fully physics-informed learning, enabling direct evaluation of the contribution of boundary constraints to predictive accuracy and data efficiency [21]. Accordingly, three surrogate models were investigated: a supervised NN, a BINN, and a PINN. Their performance was evaluated through both quantitative and qualitative comparisons under varying levels of training data.

This section describes the modeling framework and training strategy. First, the common fully connected feed-forward neural network architecture used for all surrogate models is presented. Next, a unified loss formulation for the PINN and BINN is introduced, showing how supervised, boundary-informed, and physics-informed training can be obtained within a single framework. The physical system is then described, including the reduced multiphase governing equations, boundary conditions, and input–output variables. Finally, the dataset preparation, collocation and boundary sampling strategy, optimization procedure, evaluation metrics, and the CFD reference simulations used to generate the training and validation data are presented.

2.1. Fully Connected Feed-Forward Neural Network Architecture

All surrogate models considered in this study, including the supervised NN, the BINN, and the PINN, were based on fully connected feed-forward neural network architectures. In this architecture, each neuron in a given layer was connected to every neuron in the subsequent layer, and information propagates in the forward direction from the input layer to the output layer. The neural network defines a nonlinear mapping between an input vector and an output vector according to

$$\mathbf{u} = \mathcal{N}_{\theta}(\mathbf{x}), \mathbf{x} \in X, \mathbf{u} \in U \quad (1)$$

where $\mathbf{x}$ denotes the input vector, $\mathbf{u}$ is the predicted output vector, and θ represents the set of trainable parameters, including all weights and biases. The network consists of an input layer, N_L hidden layers, and an output layer, with the overall mapping obtained through a sequence of nonlinear transformations. The output of hidden layer $l \in \{1, \dots, N_L\}$ is given by

$$\mathbf{h}^{(l)} = \sigma^{(l)} \left(\mathbf{W}^{(l)} \mathbf{h}^{(l-1)} + \mathbf{b}^{(l)} \right) \quad (2)$$

where $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ denote the weight matrix and bias vector of layer l , respectively, $\mathbf{h}^{(l-1)}$ is the output of the preceding layer with $\mathbf{h}^{(0)} = \mathbf{x}$, and $\sigma^{(l)}$ is the activation function applied elementwise. The hyperbolic tangent (tanh) activation function was selected for all hidden layers because of its smooth differentiability and compatibility with the automatic differentiation required to evaluate the PDE residuals [21]. A linear activation function ($f(x) = x$) was used in the output layer to avoid restricting the range of the predicted velocity and pressure fields [31,38]. All surrogate models were implemented and trained in MATLAB using the Deep Learning Toolbox.

2.2. Physics-Informed Neural Networks and Unified Loss Formulation

PINNs embed governing physical laws directly into the training process, enabling approximation of solution fields while reducing reliance on large labeled datasets. In this work, a PINN-based surrogate model was developed to predict the three-dimensional hydrodynamics of a gas–liquid double-impeller stirred tank. The input vector is defined as

$$x = (x, y, z, N_c) \quad (3)$$

where (x, y, z) are the spatial coordinates and N_c is the impeller rotational speed.

The objective of this study was not to develop a universal surrogate model applicable to all gas–liquid stirred tanks, but rather to establish and validate a physics-informed surrogate framework for a representative reactor configuration. Accordingly, the reactor geometry, fluid properties, and gas flow rate were held constant, while N_c , the dominant operating parameter governing gas–liquid hydrodynamics, was varied over a practical operating range. Changes in N_c directly influence the liquid circulation velocity, shear rate, turbulence intensity, and gas–liquid momentum exchange, thereby governing the overall flow structure, gas-holdup distribution, and gas-phase transport. Under these conditions, N_c provides a physically meaningful operating parameter while maintaining a compact and computationally efficient input space. The proposed framework can be readily extended to include additional operating variables, such as gas flow rate, fluid properties, or reactor geometry, in future studies. The unknown solution fields were approximated by the neural network as

$$u(x) \approx N_\theta(x) \quad (4)$$

where θ denotes the trainable network parameters. These parameters were obtained by minimizing a composite loss function that combines supervised data, boundary-condition constraints, and governing-equation residuals:

$$L(\theta) = \lambda_{PDE} L_{PDE} + \lambda_{BC} L_{BC} + \lambda_{data} L_{data} \quad (5)$$

where λ_{PDE} , λ_{BC} , and λ_{data} are the weighting coefficients associated with the PDE residuals, boundary-condition residuals, and supervised data loss, respectively [31]. These coefficients were determined through preliminary hyperparameter tuning to balance the contributions of the data, boundary, and physics loss terms, preventing any single component from dominating the optimization and promoting stable convergence with physically consistent predictions [31]. The physics-based loss enforces the governing equations at the interior collocation points:

$$L_{PDE} = \frac{1}{|X_{PDE}|} \sum_{x \in X_{PDE}} \|R(u(x))\|^2 \quad (6)$$

where $R(\cdot)$ denotes the residuals of the governing equations and X_{PDE} is the set of interior collocation points. Boundary conditions are enforced through an additional penalty term:

$$L_{BC} = \frac{1}{|X_{BC}|} \sum_{x \in X_{BC}} \|u(x) - u_{BC}(x)\|^2 \quad (7)$$

where X_{BC} denotes the set of boundary points and u_{BC} represents the prescribed boundary values. The supervised data loss penalizes discrepancies between network predictions and the labeled CFD data. Within this unified framework, the NN uses only L_{data} , the BINN uses $L_{data} + L_{BC}$, and the PINN incorporates all three loss terms, enabling direct comparison of purely data-driven, boundary-informed, and physics-informed learning [37].

Because multiphase stirred-tank flows involve coupled variables with different magnitudes and spatial characteristics, the supervised data loss was further decomposed into separate components corresponding to liquid velocity, pressure, gas holdup, and the corrective gas velocity component introduced in the reduced formulation. This decomposition prevents the optimization from being dominated by the velocity field, improving the prediction of pressure gradients and gas holdup while enhancing numerical stability and overall physical consistency. Although developed for the present stirred-tank configuration, the proposed loss formulation is readily applicable to other multiphase flow systems.

Geometric transferability was not investigated in the present study. The trained network is specific to the tank, impellers, baffles, gas-flow rate and fluid properties considered here; its validated parametric domain therefore concerns rotational speed rather than reactor geometry. A geometry-general surrogate would require geometric variables such as D/T , impeller clearance, inter-impeller spacing, baffle width and liquid height to be included in the input representation, together with CFD or experimental data spanning those geometries. Geometry-aware coordinates or signed-distance/implicit geometry representations may further facilitate such an extension.

2.3. Overview of the PINN Framework

Figure 1 illustrates the overall architecture and training workflow of the PINN, integrating the network architecture and loss formulation described above. The inputs are the spatial coordinates (x, y, z) and impeller rotational speed (N_c) , which is the sole operating parameter varied in this study. These inputs pass through a fully connected feed-forward network that predicts the liquid velocity components, pressure, gas holdup, and corrective gas velocity component. Automatic differentiation provides the spatial derivatives needed to evaluate the reduced liquid-divergence, momentum, and gas-transport residuals at interior collocation points. These residuals are combined with liquid velocity boundary losses at the tank wall, shaft, and impeller surfaces and supervised losses calculated from the CFD reference fields. The objective function is a weighted combination of the supervised, boundary, and reduced physics loss terms. Network parameters are updated using the Adam optimizer for 25,000 prescribed iterations. The individual loss components and validation error are monitored to assess whether training has stabilized. In the decision step shown in Figure 1, the No branch indicates that the prescribed iterations have not been completed and parameter updating continues; the Yes branch indicates completion of the 25,000 iterations and leads to the final trained model. Stabilization of the monitored quantities provides evidence of convergence but is not a separate numerical stopping threshold in this workflow.




Figure 1. Overview of the proposed PINN architecture and associated loss formulation.

2.4. Model Description and Reduced Multiphase Governing Equations

Previous studies have simplified the modeling of stirred bioreactors by reducing the problem from three-dimensional to two-dimensional representations because of the substantial computational and methodological challenges associated with realistic reactor geometries [29,37,39]. Although such simplifications provide valuable qualitative insight, they cannot accurately capture key hydrodynamic features of industrial stirred tanks, including impeller-driven circulation, gas dispersion, and radial and axial flow non-uniformities [40]. Consequently, this study develops a three-dimensional PINN-based surrogate model for a gas–liquid double-impeller stirred tank with a tank aspect ratio of 1.25. The model is trained using sampled CFD solution data rather than values at every location of the approximately 1.36-million-cell computational mesh. Because the spatial coordinates (x, y, z) are continuous network inputs, the trained model can be evaluated at arbitrary locations within the modeled fluid domain and is not restricted to the original CFD mesh locations. This capability represents continuous spatial interpolation of the learned solution and should not be interpreted as formal mesh-independent accuracy.

The vessel geometry is shown in Figure 2a, and its principal dimensions are summarized in Table 1. The cylindrical tank has a diameter of 0.40 m and a liquid height of 0.50 m, giving a liquid-height-to-diameter ratio of 1.25. Two pitched-blade impellers, each 0.180 m in diameter, are mounted coaxially on a central shaft. The lower impeller is positioned 0.180 m above the tank bottom, and the center-to-center spacing between the impellers is 0.180 m. Each blade is 0.038 m wide and 0.004 m thick. The tank contains four vertical baffles, each 0.040 m wide, 0.005 m thick, and 0.50 m high. A ring sparger with a diameter of 0.14 m is positioned below the lower impeller and contains 20 holes, each 0.001 m in diameter. The working fluid is an air–water system. Water is modeled as a Newtonian, incompressible liquid, and air is treated as having constant density under the operating conditions considered. In the CFD model, air enters through the ring sparger at a velocity inlet corresponding to the prescribed gas flow rate, with a gas volume fraction of $\alpha_g = 1.0$ at the inlet. A degassing condition at the upper liquid surface permits gas outflow while preventing liquid loss. The solid-surface boundary conditions used in the CFD model are specified for the tank walls, baffles, impellers, and shaft in Section 2.9.

Steady-state training data were generated from CFD simulations using the multiple-reference-frame (MRF) approach, which is widely adopted for modeling rotating impellers in aerated stirred tanks [41–44]. The CFD velocity components used as supervised targets were exported as absolute velocities in a common Cartesian coordinate system for the stationary and rotating regions. The PINN predicts velocity in that same frame, so no relative-to-absolute-velocity transformation was applied to the supervised targets. The reference CFD solution nevertheless accounts for impeller rotation through its multiple-reference-frame formulation. The PINN's reduced momentum residual does not reproduce the complete MRF momentum treatment or the CFD turbulence and interphase momentum-transfer terms. Consequently, agreement with the exported CFD fields must be attributed to the combined supervised data, boundary constraints, and reduced physics regularization; the PINN residual is not equivalent to the CFD momentum equation.

The surrogate model directly predicts the liquid velocity components in the x -, y -, and z -directions, together with the pressure and gas-holdup fields throughout the computational domain. The gas velocity is then reconstructed as the sum of the predicted liquid velocity, a slip velocity computed from the predicted gas holdup, and a learned corrective-velocity component that accounts for unresolved multiphase interactions. The reactor geometry is shown in Figure 2, and the corresponding dimensions and fluid properties are summarized in Table 1.




Figure 2. (a) Three-dimensional geometry of the double-impeller stirred tank used for CFD simulations and surrogate model development, with the regions of interest (ROIs) highlighted in shaded bands around each impeller; (b) key geometrical parameters and dimensions.

Table 1. Geometrical dimensions of the double-impeller stirred-tank configuration.

Quantity	Dimension (m)
D_{impeller}	0.180
W_{Baffle}	0.040
t_{Baffle}	0.005
D_{tank}	0.400
H_{Tank}	0.500

The benchmark case considered in this study is the steady-state operation of a three-dimensional gas–liquid double-impeller stirred tank. The liquid phase is treated as a Newtonian, incompressible fluid with constant density and viscosity, while the gas phase is represented through a continuous Eulerian volume-fraction field describing the local gas holdup. These assumptions are consistent with commonly adopted modeling approaches for aerated stirred tanks and provide a practical basis for representing reactor-scale gas–liquid hydrodynamics [16,45,46].

The reference CFD simulations employ a more complete Euler–Euler multiphase formulation, including turbulence and interphase closures, as described in Section 2.9. In contrast, reproducing the complete CFD formulation within the PINN would require additional phase-specific momentum equations, turbulence transport equations, and interphase closure relations, considerably increasing the complexity of the physics-informed optimization. Therefore, the PINN employs a reduced multiphase formulation as a set of physics-based regularization constraints, rather than attempting to reproduce the complete Euler–Euler/RANS equations solved in the reference CFD model. In this reduced formulation, the gas-phase momentum equation is not solved explicitly. Instead, gas transport is represented using the predicted liquid velocity, an algebraic slip relation, and a learned corrective velocity. The effects of turbulence and unresolved interphase interactions are therefore incorporated primarily through the supervised CFD fields and, for gas transport, through the corrective component of the reconstructed gas velocity.

Accordingly, the proposed framework is intended to predict the reactor-scale mean hydrodynamic fields, including liquid velocity, pressure, and gas holdup, rather than explicitly resolve bubble-scale dynamics, turbulence quantities, or individual interphase forces. The surrogate predictions are evaluated against CFD solution fields from a model

validated using experimental measurements of gas holdup and gassed power consumption, as described in Section 2.9 [47]. For an incompressible liquid of constant density, the full steady Euler–Euler liquid-phase continuity equation can be written as $\nabla \cdot (\alpha_l u_l) = 0$, where $\alpha_l = 1 - \alpha_g$ is the liquid volume fraction. In the present PINN, this relation is simplified to the following liquid velocity incompressibility constraint:

$$R_{\text{cont.}} = \nabla \cdot u_l = 0 \quad (8)$$

where u_l denotes the liquid velocity field. Equation (8) is a reduced constraint, not the complete liquid-phase continuity equation: when liquid volume fraction varies spatially, $\nabla \cdot (\alpha_l u_l) = 0$ does not generally imply $\nabla \cdot u_l = 0$. The momentum contribution to the PINN loss likewise uses a reduced residual with gas-holdup-dependent effective properties:

$$R_{\text{mom.}} = \rho_{\text{eff}}(\alpha_g)(u_l \cdot \nabla)u_l + \nabla p - \mu_{\text{eff}}(\alpha_g)\nabla^2 u_l - \rho_{\text{eff}}(\alpha_g)g = 0 \quad (9)$$

where g is the gravitational acceleration, and the effective density and viscosity are defined as

$$\rho_{\text{eff}}(\alpha_g) = (1 - \alpha_g)\rho_l + \alpha_g\rho_g \quad (10)$$

$$\mu_{\text{eff}}(\alpha_g) = (1 - \alpha_g)\mu_l + \alpha_g\mu_g \quad (11)$$

The effective properties are phase-fraction-weighted combinations of the individual phase properties [48,49]. Their use in Equation (9) does not make that residual the complete liquid-phase or mixture momentum equation of the Euler–Euler CFD model. Equations (8) and (9) serve as physics-guided regularizers. Equation (8) omits the effect of spatially varying liquid volume fraction on phase continuity, while Equation (9) does not explicitly include the CFD turbulence and interphase momentum-transfer terms. The reference CFD fields have not been shown to satisfy these reduced residuals exactly; the residuals should therefore not be presented as an independently verified reproduction of the CFD equations.

Equation (9) is formulated in the inertial frame. The reference CFD simulations used an absolute-velocity multiple-reference-frame (MRF) formulation, and the exported absolute-velocity fields were used as supervised targets without a relative-to-absolute-velocity transformation. The reduced PINN residual does not reproduce all terms of the MRF-based CFD momentum formulation. Information about impeller rotation enters the surrogate through the supervised CFD fields and prescribed rotational boundary conditions.

The reduced momentum residual also omits explicit turbulent-viscosity or Reynolds-stress contributions and does not solve separate turbulence-transport equations. The reference CFD fields were obtained using the standard k – ϵ turbulence model; turbulence therefore influences the surrogate through its supervised mean-flow targets. No additional eddy-viscosity closure or learned momentum-source term was introduced. The PINN is consequently a CFD-informed surrogate with reduced physics regularization, rather than an independent solver of the complete MRF/RANS/Euler–Euler system. Adding an eddy-viscosity contribution or a learned momentum-discrepancy term could be investigated in future work. The gas velocity is reconstructed as:

$$u_g = u_l + u_{\text{slip}}(\alpha_g) + \Delta u_g \quad (12)$$

where the gas velocity, u_g , is decomposed into u_l , representing the liquid convective velocity; u_{slip} , representing the prescribed hindered-rise slip contribution; and Δu_g , representing the learned corrective velocity. The slip velocity, u_{slip} , is given by:

$$u_{slip}(\alpha_g) = \begin{bmatrix} 0 \\ u_t(1 - \alpha_g)^{n_h} \\ 0 \end{bmatrix} \quad (13)$$

where u_t is the terminal-rise velocity scale and n_h is the hindrance exponent characterize the algebraic slip relation. The terminal-rise velocity scale and hindrance exponent define the prescribed slip relation. The learned corrective velocity is a data-informed adjustment to that relation. It cannot be uniquely attributed to turbulence, interphase momentum transfer, or another individual mechanism omitted from the reduced formulation. The global field-norm ratios reported in Section 3.2.5 quantify the relative magnitude of the learned corrective velocity over the evaluated domain. They do not characterize its spatial distribution or establish that the correction is small in every flow region. Because a separately trained model without the corrective term was not evaluated, its effect on predictive accuracy has not been isolated. Gas transport is constrained through the steady gas volume-fraction residual:

$$R_g = \nabla \cdot (\alpha_g \mathbf{u}_g) = 0 \quad (14)$$

This residual uses the reconstructed gas velocity from Equation (12). The CFD gas-inlet and outlet conditions are represented in the supervised CFD fields, but separate gas velocity or gas-holdup boundary losses are not imposed at the sparger inlet or upper outlet. Consequently, the PINN residual and boundary losses do not independently enforce the prescribed global inlet-to-outlet gas flux.

Together, Equations (8), (9), and (14) form the physics residual vector used for PINN regularization:

$$\mathbf{R} = [R_{\text{cont}}, R_{\text{mom}}, R_g]^T \quad (15)$$

These are reduced physics-guided constraints, not the complete governing equations of the reference Euler–Euler/RANS model. The network outputs are defined as:

$$\mathbf{u} = (u_{l,x}^*, u_{l,y}^*, u_{l,z}^*, p^*, \alpha_{\text{raw}}, \Delta u_{g,x}^*, \Delta u_{g,y}^*, \Delta u_{g,z}^*)^T \quad (16)$$

They comprise nondimensional liquid velocity components, nondimensional pressure, an unconstrained gas-holdup variable, and nondimensional corrective gas velocity components. The gas-holdup output is mapped to the interval between zero and one through a sigmoid transformation:

$$\alpha_g = \sigma(\alpha_{\text{raw}}) \in [0, 1] \quad (17)$$

To represent the dependence on operating conditions, the network input is defined as:

$$\mathbf{X} = [N_c^*, x^*, y^*, z^*]^T \quad (18)$$

The nondimensional impeller speed, $N_c^* = N_c / N_{c,\text{ref}}$, is defined using a reference rotational speed. Including it as an input allows one network to represent the speeds covered by the training data. Predictions at unseen intermediate speeds are interpolation within the trained 200–500 rpm range, not evidence of extrapolation beyond that range.

2.5. Boundary Conditions and Pressure Anchoring

Boundary conditions were incorporated through a boundary-loss term applied at selected boundary points. No-slip boundary conditions were prescribed on the stationary

tank wall and bottom surface. The rotating motion of the impellers and shaft was represented using a moving-wall condition prescribed on the impeller blade and shaft surfaces:

$$\mathbf{u}_l = \boldsymbol{\omega} \times \mathbf{r} \quad (19)$$

where $\boldsymbol{\omega} = [0, \omega, 0]^T$ is the angular velocity vector aligned with the rotation axis, $\omega = 2\pi N_c/60$, and $\mathbf{r}$ is the position vector relative to the rotation axis.

No dedicated gas-phase boundary-loss terms were imposed in the present PINN formulation. In particular, gas velocity and gas holdup at the sparger inlet and upper gas-outflow boundary were not prescribed independently during training. Their boundary behavior was learned from the supervised CFD fields, while gas transport within the domain was additionally regularized through the gas volume-fraction conservation residual in Equation (14). Therefore, the gas-phase boundary conditions of the reference CFD model are represented implicitly rather than enforced as separate PINN constraints.

In incompressible flow, the pressure field is determined only up to an arbitrary additive constant because the momentum residual depends on pressure gradients rather than absolute pressure. A two-step gauge-fixing strategy was therefore adopted to define a unique pressure reference and improve numerical stability during training. First, the CFD pressure field was shifted to zero spatial mean for each operating condition during data preprocessing, preserving the physically relevant pressure gradients while removing the arbitrary constant-pressure mode. Second, a small number of pressure-anchor points were included in the loss function by enforcing

$$p(x_p, k) = 0, \quad k = 1, \dots, N_p \quad (20)$$

These reference constraints prevent arbitrary pressure offsets and pressure drift during training, thereby improving the numerical stability of the surrogate. This treatment is consistent with standard pressure-gauge approaches for incompressible flow and their application in physics-informed fluid-flow models [38,50].

2.6. Dataset Preparation and Region-Aware Sampling of Labeled Data

Experimentally validated three-dimensional CFD simulations were used to generate the labeled data for training, validation, and testing of the surrogate models. The dataset consists of the spatial coordinates together with the corresponding liquid velocity components, gas velocity components, pressure, and gas holdup. To limit the number of supervised training samples while preserving information in dynamically important regions, a region-aware sampling strategy was adopted, consistent with strategic spatial sampling approaches for neural network training in complex flows [21,51]. The final supervised dataset contained 200,000 labeled CFD points sampled across the four training conditions, with 50,000 points retained at each impeller speed of $N_c = 200, 300, 400$, and 500 rpm.

Two axial regions extending ± 0.04 m from the lower and upper impeller planes were defined as regions of interest (ROIs). This distance was selected to encompass the primary impeller discharge regions, where the highest velocity gradients, strongest turbulence, and most intense gas–liquid interactions occur. As these regions dominate the reactor hydrodynamics and strongly influence the global circulation pattern, all labeled data within the ROIs were retained to preserve the spatial resolution required to accurately learn the impeller-driven flow structures [16,50]. The region between the two impellers was allocated approximately 25% of the labeled samples to capture the inter-impeller circulation and gas dispersion that strongly influence the overall mixing performance. The remaining labeled samples were distributed uniformly throughout the rest of the tank volume, excluding the ROIs, to ensure representative coverage of the global flow field. The 200,000-point

budget refers to the total number of labeled CFD samples before partitioning; it comprises 160,000 training, 20,000 validation, and 20,000 test points. The random pointwise partition was used to evaluate prediction at spatial coordinates not included in the supervised training subset. Because the samples were extracted from continuous CFD fields, spatially neighboring points may occur in different subsets; consequently, this test should not be interpreted as a fully spatially independent validation. Rather, it evaluates pointwise spatial interpolation within the CFD-derived solution domain. A more stringent assessment based on spatially blocked or plane-wise holdout regions would be required to evaluate generalization to spatial regions completely excluded from the labeled training data.

2.7. Collocation Sampling and Spatial Weighting for Physics Enforcement

Physics-based residuals were enforced at interior collocation points sampled throughout the fluid domain. To improve the resolution of the dominant flow structures, a spatially biased sampling strategy was adopted instead of uniform sampling [30,52]. During the first 6500 training iterations, 15,000 interior collocation points were sampled at each iteration, after which the collocation batch size was increased to 30,000 points per iteration for the remaining 18,500 iterations. The collocation points were resampled at every training iteration. Specifically, 40% of each collocation batch was sampled from an axial band defined by $-0.02 \leq z \leq 0.02$ m, corresponding to a total width of 0.04 m and centered on the tank mid-plane. This band was selected because it intersects the primary circulation region, where the impeller discharge flows interact, gas plumes merge, and strong flow recirculation occurs. The remaining 60% of the collocation points were sampled throughout the full fluid domain to ensure representative coverage of the global flow field, including wall-affected regions. For each collocation point, the impeller speed was randomly assigned from the four training conditions ($N_c = 200, 300, 400$, and 500 rpm), such that the reduced physics-based residuals were evaluated across the investigated operating range.

In addition to biased collocation sampling, the PDE residuals were spatially weighted to increase the contribution of dynamically important flow regions to the physics loss [20]. The dimensions of the weighted regions were selected based on preliminary inspection of the CFD flow fields and sensitivity analyses to encompass the dominant impeller discharge and circulation zones while limiting unnecessary weighting of the surrounding bulk flow. Collocation points within the impeller regions (within ± 0.01 m radially and ± 0.019 m axially of each impeller) were assigned a weight of 5.0 because these regions contain the strongest velocity gradients and gas–liquid interactions, making them the most challenging locations for satisfying the reduced physics-based constraints. Uniform weighting tends to underemphasize these localized features in favor of the smoother bulk flow. The selected weight was determined through preliminary sensitivity analysis: lower weights underpredicted the impeller-driven jet structures, whereas higher weights reduced training stability. Points in the inter-impeller region ($0.22 \leq y \leq 0.32$ m) were assigned a moderate weight of 2.0 to better resolve the inter-impeller circulation and gas plume interactions. Near-wall ($r \geq 0.95 \times \text{tank radius}$) and near-bottom ($y \leq 0.05 \times \text{liquid height}$) regions were assigned a reduced weight of 0.3 because their thin boundary layers can otherwise dominate the loss function despite having limited influence on the overall reactor hydrodynamics. All remaining collocation points were assigned a baseline weight of 1.0. All remaining collocation points were assigned a baseline weight of 1.0. Overall, the combined sampling and spatial-weighting strategy concentrates the reduced physics-based constraints in the most dynamically important regions while maintaining adequate coverage of the computational domain and stable optimization.

To assess whether the manually specified collocation strategy could be improved through adaptive point redistribution, an additional residual-driven adaptive collocation

strategy was evaluated. Candidate interior points were periodically evaluated using the current reduced physics residuals, and points associated with larger residual magnitudes were assigned a higher probability of selection in subsequent collocation batches. Consequently, the spatial distribution of the collocation points evolved during training toward regions in which the reduced physics-based constraints were more difficult to satisfy, rather than relying exclusively on the predefined spatial regions used in the baseline strategy. For a controlled comparison, the network architecture, supervised CFD dataset, total collocation-point budget, optimizer, training duration, and loss weights were maintained identical to those of the baseline model. The resulting predictive performance was compared with that obtained using the baseline heuristic collocation strategy, as discussed quantitatively in Section 3.2.

2.8. Training Schedule, Optimization, and Stability Measures

Training was performed by minimizing a composite loss consisting of supervised data, boundary-condition, and physics-based terms. The supervised loss penalized errors in the predicted liquid velocity, pressure, gas holdup, and corrective gas velocity component (Δu_g). Including Δu_g in the supervised loss constrained the learned correction to remain consistent with the CFD-predicted gas velocity while allowing the prescribed slip relation to capture the dominant contribution to gas transport. Boundary losses were applied only to the liquid velocity because the solid boundaries directly influence the liquid phase through no-slip and moving-wall conditions. Since the gas velocity is reconstructed from the liquid velocity and the holdup-dependent slip relation, no separate boundary constraints were imposed on the gas velocity. A pressure-anchor loss was included to remove the pressure null space associated with the arbitrary constant offset and improve numerical stability during training. In addition, a smoothing regularization term was applied to the gas holdup field (α_g) to suppress nonphysical high-frequency oscillations. A penalty term was also imposed on Δu_g to prevent the learned correction from dominating the prescribed slip relation, thereby restricting it to a small refinement of the gas-phase velocity. To evaluate the sensitivity of the learned corrective velocity to the assumptions associated with its formulation, a local sensitivity analysis was performed by independently perturbing the corrective-velocity regularization weight, $\lambda_{\Delta u}$, the terminal rise velocity, u_t , and the hindrance exponent, n_h , by $\pm 20\%$ relative to their baseline values while maintaining the remaining model and training parameters unchanged. The relative magnitude of the learned correction was quantified using $R_{\text{slip}} = \|\Delta u_g\|_2 / \|u_{\text{slip}}\|_2$ and $R_l = \|\Delta u_g\|_2 / \|u_l\|_2$. The corresponding effects on gas velocity and gas-holdup prediction errors were also evaluated. The results of this sensitivity analysis are presented in Section 3.2. To improve training stability, the network was initially trained using the supervised data loss alone. The PDE, gas-continuity, boundary-condition, and pressure-anchor loss weights were then increased from zero to their target values over 8000 iterations. For the remaining iterations, these weights were held fixed. This schedule gradually introduced the additional constraints while retaining the supervised CFD targets throughout training.

Optimization was performed using the Adam (Adaptive Moment Estimation) algorithm, a gradient-based optimization method that adaptively updates the learning rate for each network parameter using estimates of the first and second moments of the gradients [53]. The learning rate was set to $\eta = 5 \times 10^{-4}$, with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The learning rate was selected through preliminary hyperparameter tuning to balance convergence speed and numerical stability. Values of $\eta \leq 10^{-4}$ resulted in slow convergence, whereas $\eta \geq 10^{-3}$ frequently caused divergence following activation of the PDE residuals. The remaining Adam parameters were retained at their standard values [53].

The learning rate was selected through preliminary hyperparameter tuning to balance convergence speed and numerical stability. Lower values ($\leq 10^{-4}$) led to slow convergence, whereas higher values ($\geq 10^{-3}$) frequently caused divergence after activation of the PDE residuals. The remaining Adam hyperparameters were kept at their standard values, which have been shown to perform well in PINN-based fluid-flow applications [20].

Separate mini-batches were used for the supervised data, interior collocation points, and boundary points to maintain balanced gradient contributions while remaining computationally efficient [20]. Global gradient-norm clipping with a threshold of 5 was applied to prevent occasional large parameter updates, particularly during the transition to physics-informed training [54]. Model performance was monitored using the validation dataset throughout training, and the final model was evaluated on the independent test dataset.

The loss weights were selected through preliminary tuning and are treated as training hyperparameters rather than universal model constants. To assess sensitivity to their selection, the reduced momentum and gas-continuity weights were independently varied by $\pm 50\%$ from their baseline values of 0.02 and 0.03, respectively, while the remaining training settings were held fixed. The resulting prediction errors are discussed in Section 3.2.

2.9. CFD Reference Simulations and Validation Strategy

Three-dimensional, steady-state CFD simulations of gas–liquid flow in the double-impeller stirred tank were performed using ANSYS Fluent 2025 R2 to generate reference fields for surrogate-model training and evaluation. The simulations used a pressure-based solver with SIMPLE pressure–velocity coupling. The momentum equations were discretized using a second-order upwind scheme, and the gas volume-fraction equation was discretized using the QUICK scheme. Impeller rotation was represented using the multiple-reference-frame (MRF) approach, with rotating zones around the impellers and a stationary tank region. The simulations used the absolute-velocity formulation; absolute-velocity fields exported from Fluent were therefore used as supervised targets without a relative-to-absolute-velocity transformation.

The gas–liquid flow was modeled using an Euler–Euler formulation with a representative bubble diameter of 0.003 m. Turbulence was modeled using the standard k – ϵ model, and interphase momentum exchange was represented using the Schiller–Naumann drag correlation. Lift, virtual-mass, wall-lubrication, and turbulent-dispersion forces were not explicitly included, and bubble breakup and coalescence were not resolved using a population balance model. The resulting turbulence-closed mean velocity, pressure, and gas-holdup fields were used as supervised targets for the NN, BINN, and PINN. None of the surrogate models solves transport equations for turbulent kinetic energy (k) or its dissipation rate (ϵ). The PINN’s reduced momentum residual also omits explicit turbulent-viscosity and Reynolds-stress terms. Thus, turbulence influences the surrogate through the supervised CFD fields, while the reduced residual serves as a physics-guided constraint rather than reproducing the complete CFD momentum equation (Section 2.4).

Reference simulations were conducted at a constant gas flow rate of 20 L min^{−1} and impeller speeds of 200, 300, 400, and 500 rpm. Air entered through a 0.14 m diameter ring sparger below the lower impeller. A velocity-inlet condition corresponding to the prescribed gas flow rate was applied at the sparger openings, with a gas volume fraction of 1.0 at the inlet. A degassing condition at the upper liquid surface permitted gas outflow while preventing liquid loss. No-slip conditions were applied to the stationary tank walls and baffles; the impeller and shaft surfaces were treated according to their rotation. The liquid height was 0.50 m. Each speed was initialized using a converged solution from an adjacent operating condition.

Grid independence was assessed using meshes containing 741,554, 1,363,732, and 2,671,865 cells. The medium mesh was used for subsequent simulations because its volume-averaged gas-holdup prediction differed from that of the fine mesh by less than 3.8%. The simulations were run on a high-performance computing (HPC) system using 32 CPU cores until the residuals and monitored gas holdup and gassed power consumption reached stable values. Each operating condition required approximately 100–140 h of wall-clock time. The converged solution fields were exported at CFD mesh locations and sampled to construct the neural network datasets.

The CFD reference model was compared with experimental measurements obtained for the same vessel and impeller geometry at 20 L min^{−1} and impeller speeds of 400 and 500 rpm. The measured quantities were overall gas holdup, local gas holdup at four electrical resistance tomography planes, and total gassed power consumption. Table 2 reports the experimental and CFD values for each condition. The displayed local gas-holdup deviations range from 4.19% to 8.70%. These comparisons support the use of the CFD fields as reference data at the two experimentally assessed speeds; they do not constitute experimental validation at every training or interpolation speed.

Table 2. Validation of CFD-predicted overall and local gas holdup and total gassed power consumption against experimental measurements at impeller speeds of 400 and 500 rpm.

N_c (rpm)	Q_g (L/min)	Total Gassed Power Consumption, P_g (W)		Error (%)	Overall Gas Holdup, ε_g (-)		Error (%)	ERT plane	Local Gas Holdup, ε_g (-)		Error (%)
		EXP	CFD		EXP	CFD			EXP	CFD	
400	20	17.12	17.84	4.21	0.102	0.108	5.88	#1	0.083	0.089	7.23
								#2	0.104	0.110	5.77
								#3	0.125	0.133	6.4
								#4	0.115	0.122	6.09
500	20	33.4	35.08	5.03	0.134	0.141	5.22	#1	0.115	0.125	8.7
								#2	0.131	0.139	6.11
								#3	0.165	0.173	4.58
								#4	0.143	0.148	3.58

The labeled CFD points were randomly partitioned into training, validation, and test subsets. Test points were excluded from the training and validation subsets, but neighboring points from the same continuous CFD field could occur in different subsets. The resulting test therefore assesses pointwise spatial interpolation rather than generalization to an entirely withheld spatial region. The PINN was also evaluated against CFD results at 350 and 450 rpm, which were excluded from training but lie within the 200–500 rpm training range. Because the NN and BINN were not evaluated at these intermediate speeds, those results assess PINN speed interpolation without establishing an interpolation advantage over the baseline models.

2.10. Error Metrics

Model accuracy was quantified using normalized pointwise errors for the liquid velocity, pressure, gas holdup, and reconstructed gas velocity. For a scalar field ϕ , the normalized L_1 , normalized L_2 , and root-mean-square error (RMSE) are defined as

$$L_1 = \frac{1}{N} \sum_{i=1}^N \frac{|\phi_i^{\text{pred}} - \phi_i^{\text{CFD}}|}{\phi_{\text{ref}}} \quad (21)$$

$$L_2 = \left(\frac{1}{N} \sum_{i=1}^N \left(\frac{\phi_i^{\text{pred}} - \phi_i^{\text{CFD}}}{\phi_{\text{ref}}} \right)^2 \right)^{1/2} \quad (22)$$

$$\text{RMSE} = \left(\frac{1}{N} \sum_{i=1}^N (\phi_i^{\text{pred}} - \phi_i^{\text{CFD}})^2 \right)^{1/2} \quad (23)$$

where ϕ denotes the predicted variable and ϕ_{CFD} is the corresponding CFD reference field. For the vector-valued liquid and gas velocity fields, each metric was first evaluated separately for the three Cartesian velocity components (u , v , and w), and the reported velocity metric was obtained by taking the arithmetic mean of the three component-wise values. Thus, the three velocity components contributed equally to the reported aggregate error, irrespective of differences in their individual magnitudes. The same component-wise equal-weighting procedure was used for the normalized L_1 and L_2 errors, RMSE, and R^2 . Velocity components were scaled using the impeller tip speed, while pressure and gas holdup were scaled using the standard deviation of their respective CFD reference fields where scaling was required by the metric definition.

2.11. Model Architecture and Training Configuration

The main network architecture, optimization settings, loss weights, and multiphase slip-model parameters used throughout this study are summarized in Table 3. A fully connected multilayer perceptron (MLP) with six hidden layers and 128 neurons per layer was adopted to provide sufficient representational capacity while maintaining stable training. The hyperbolic tangent activation function was used in all hidden layers, and optimization was performed using the Adam optimizer with a constant learning rate of 5×10^{-4} . A two-stage training schedule consisting of a supervised warm-up phase followed by gradual activation of the physics-based loss terms was employed to improve convergence. The loss weights were selected through preliminary sensitivity analysis to balance supervised learning and physics enforcement while maintaining comparable gradient contributions during optimization. Training was performed for 25,000 iterations, after which both the validation loss and PDE residuals had converged with no significant further improvement. To assess the sensitivity of the trained PINN to stochastic initialization, five independent training runs were performed using different random seeds while maintaining identical model and training settings. A local sensitivity analysis was also conducted by varying the reduced physics and gas-continuity loss weights by $\pm 50\%$ relative to their baseline values while keeping the remaining hyperparameters unchanged.

Table 3. Summary of the main model architecture and training parameters.

Category	Hyperparameter	Value
Network architecture	Network type	Fully connected feed-forward MLP
	Hidden layers	6
	Neurons per hidden layer	128
	Activation function	Tanh
	Input/output dimensions	4/8
Training setup	Optimizer	Adam
	Learning rate	5×10^{-4}
	Training iterations	25,000
	Gradient clipping	Global norm = 5
Mini-batch configuration	Supervised data	5000 points/iteration
	Interior collocation, iterations 1–6500	15,000 points/iteration
	Interior collocation, iterations 6501–25,000	30,000 points/iteration

Table 3. Cont.

Category	Hyperparameter	Value
Loss formulation	Data loss weight, λ_d	10
	Reduced momentum/continuity loss weight, λ_p	0.02
	Gas-continuity loss weight, λ_g	0.03
	Boundary-condition loss weight, λ_b	1
	Pressure-anchor loss weight, λ_{pa}	0.7
	Gas-holdup smoothing weight, $\lambda_{\alpha,smooth}$	1×10^{-6}
	Corrective-velocity penalty weight, $\lambda_{\Delta u}$	1×10^{-4}
Training schedule	Data-only warm-up	2500 iterations
	Physics/loss ramp-up	8000 iterations
Multiphase modeling	Slip-velocity model	Algebraic hindered-rise
	Terminal rise velocity, u_t	0.25 m s^{-1}
	Hindrance exponent, n_h	2
Data handling	Total labeled CFD dataset	200,000 points
	Train/validation/test split	80/10/10%

3. Results and Discussion

This section presents a comprehensive evaluation of the proposed framework. The analysis is organized to examine (i) the influence of labeled-data availability on predictive accuracy, (ii) the comparative performance of supervised, boundary-informed, and physics-informed neural networks, and (iii) the ability of the PINN to interpolate at intermediate operating conditions not included during training. Model performance was assessed using both quantitative error metrics and qualitative flow-field comparisons to evaluate predictive accuracy, physical consistency, and robustness across varying impeller speeds.

3.1. Impact of Labeled Dataset Size on Prediction Accuracy

To quantify the amount of CFD data required to train an accurate surrogate model, a label-budget study was performed in which the number of supervised training points was systematically varied while keeping the network architecture, physics formulation, and training hyperparameters fixed. The proposed PINN was trained across multiple operating conditions corresponding to impeller rotational speeds of $N_c = 200, 300, 400$, and 500 rpm. For each case, the supervised dataset size ranged from 5000 to 400,000 labeled points to evaluate the effect of data availability on predictive accuracy. Prediction accuracy was assessed using the normalized L_1 error computed for the liquid velocity, gas velocity, and pressure fields. Results at $N_c = 400$ are shown in Figure 3 for both the purely supervised NN and the PINN. Increasing the number of labeled points consistently reduces prediction error for all quantities; however, the PINN achieves substantially lower errors than the purely supervised NN across all label budgets, demonstrating improved data efficiency. For example, the normalized L_1 error of the axial gas velocity component decreases from approximately 49.5% with 5000 labeled points to 15.4% at 100,000 points. Further increasing the dataset to 200,000 points reduces the error to approximately 9.6%, with diminishing reductions to 7.4% and 6.1% at 300,000 and 400,000 points, respectively. A similar trend is observed for pressure prediction, where the normalized L_1 error decreases from 35.6% at 5000 labeled points to 11.7% at 100,000 points and falls below 5.0% at 200,000 labeled points.




Figure 3. Variation in normalized L_1 error (%) with the number of labeled training data for gas velocity and pressure fields at $N_c = 400$ rpm.

These results indicate that while increasing the number of labeled points yields substantial accuracy gains at low to moderate label budgets, the rate of improvement diminishes beyond approximately 200,000–300,000 labeled points, highlighting a saturation trend in predictive accuracy. At very low label budgets, the largest errors are concentrated in regions characterized by strong gradients and intense multiphase interactions, particularly near the impeller discharge zones and the shaft, where sparse supervision limits the network's ability to resolve localized flow structures. As the label budget increases, these dynamically important regions become better constrained, leading to marked reductions in velocity and pressure errors and a more stable reconstruction of the gas velocity field. Based on the observed trade-off between accuracy improvement and computational cost, a labeled dataset size of 200,000 CFD points was selected for the remainder of this study. This label budget provides near-converged accuracy for all predicted fields while avoiding the increased training time and limited performance gains observed at higher data volumes. Using this fixed data budget, the following sections compare the predictive performance of the supervised NN, BINN, and PINN formulations.

3.2. Comparison of Surrogate Modeling Strategies (NN, BINN, and PINN)

The NN, BINN, and PINN models were evaluated by comparing their predictions of the liquid velocity, gas velocity, pressure fields and gas holdup in the aerated double-impeller stirred tank. This comparison demonstrates how progressively incorporating boundary conditions and reduced physics-based constraints improves agreement with the CFD reference fields. Although the models were trained and evaluated across a range of impeller speeds ($N_c = 200, 300, 400$, and 500 rpm), the discussion focuses on the representative cases of $N_c = 400$ and 500 rpm, where the flow is more complex and gas–liquid interactions are more pronounced. Quantitative error metrics for the four predicted hydrodynamic quantities are summarized in Table 4, and qualitative comparisons of the predicted flow fields are presented in Figures 4–7.

Table 4. Quantitative performance metrics for the purely supervised NN, BINN, and PINN in predicting gas velocity (u_g), liquid velocity (u_l), pressure (p) fields, and gas holdup (α_g) at $N_c = 400$ and $N_c = 500$ rpm.

		NN Supervised				BINN				PINN			
		RMSE *	R^2 (-)	L_1 (%)	L_2 (%)	RMSE *	R^2 (-)	L_1 (%)	L_2 (%)	RMSE *	R^2 (-)	L_1 (%)	L_2 (%)
u_g	$N_c = 400$	0.2644	0.7712	74.75	57.03	0.1158	0.8569	37.54	29.99	0.0829	0.9573	8.98	8.48
	$N_c = 500$	0.2784	0.6728	58.50	43.87	0.1486	0.8352	34.62	27.42	0.1099	0.9486	8.88	8.40
u_l	$N_c = 400$	0.1174	0.8352	56.10	25.33	0.1159	0.9169	26.48	21.99	0.0638	0.9697	9.07	8.57
	$N_c = 500$	0.147	0.8366	53.99	23.17	0.1481	0.9056	24.46	20.35	0.0689	0.9586	8.88	8.40
p	$N_c = 400$	209.2	0.8775	18.38	15.80	208.1	0.9353	11.23	15.71	71.695	0.9831	3.09	3.75
	$N_c = 500$	209.2	0.8576	18.62	15.27	213.1	0.9258	12.81	15.55	88.27	0.9802	3.46	4.46
α_g	$N_c = 400$	0.0153	0.8118	57.90	26.65	0.0091	0.8912	28.12	23.13	0.0042	0.93285	9.18	8.67
	$N_c = 500$	0.0214	0.8048	56.26	24.61	0.0115	0.8712	26.03	21.54	0.0044	0.92217	8.99	8.50

* RMSE is reported in physical units (m/s for velocity and Pa for pressure); RMSE for gas holdup is dimensionless.

In addition to comparing the three surrogate formulations, we assessed the PINN's sensitivity to random initialization and loss-weight selection. For the initialization assessment, the PINN was trained five times with different random seeds while keeping the dataset, architecture, and other training settings fixed. The normalized L_1 errors for the liquid velocity, gas velocity, pressure, and gas holdup were $8.96 \pm 0.31\%$, $8.91 \pm 0.34\%$, $3.30 \pm 0.18\%$, and $8.47 \pm 0.39\%$ respectively.

The sensitivity of the PINN to the selected physics-loss weights was evaluated by independently varying the reduced momentum weight from 0.01 to 0.03 and the gas-continuity weight from 0.015 to 0.045. Across these perturbations, the normalized errors ranged from approximately 8.7–9.4% for liquid velocity, 8.6–9.5% for gas velocity, 3.1–3.7% for pressure, and 8.1–9.0% for gas holdup. The prediction errors changed moderately within the investigated ranges, with gas velocity showing the greatest variation. This local analysis indicates that the reported accuracy is not strongly dependent on the exact baseline weights within the tested ranges; it does not establish robustness to all possible weight selections or random initializations.

The influence of the collocation strategy was further evaluated by comparing the baseline heuristic spatial weighting with residual-driven adaptive collocation while maintaining the network architecture, supervised dataset, total collocation-point budget, optimizer, training duration, and loss weights unchanged. At $N_c = 400$ rpm, adaptive collocation reduced the normalized L_1 errors for liquid velocity, gas velocity, pressure, and gas holdup from 9.07%, 8.98%, 3.09%, and 9.18% to 8.62%, 8.45%, 2.96%, and 8.67%, respectively. At $N_c = 500$ rpm, the corresponding errors decreased from 8.88%, 8.88%, 3.46%, and 8.99% to 8.47%, 8.32%, 3.31%, and 8.51%, respectively. Overall, residual-driven adaptive collocation resulted in relative reductions of approximately 4–6% in the global prediction errors. These modest but consistent improvements indicate that adaptive redistribution of the collocation points provides additional refinement, while the baseline heuristic strategy already provides effective coverage of the principal flow regions.

3.2.1. Corrective Gas Velocity Analysis

The contribution of the learned corrective gas velocity, Δu_g , was assessed relative to the prescribed slip velocity and predicted liquid velocity using field-norm ratios. In the baseline case, $\|\Delta u_g\| / \|u_{\text{slip}}\| = 0.1460$ and $\|\Delta u_g\| / \|u_l\| = 0.0438$. Thus, the correction represents approximately 15% of the slip-field norm and 4% of the liquid velocity field norm over the evaluated domain. These global ratios characterize the overall contribution of the correction but do not establish that it is small at every spatial location.

Table 5 summarizes a local sensitivity analysis in which the corrective-velocity regularization weight, terminal rise velocity, and hindrance exponent were independently varied by $\pm 20\%$. Across these cases, the corrective-to-slip-norm ratio ranged from 0.1147 to 0.1877, and the corrective-to-liquid-norm ratio ranged from 0.0376 to 0.0522. Gas velocity RMSE ranged from 0.0369 to 0.0421 m/s, compared with 0.0379 m/s for the baseline; the largest relative increase was 11.08%. Gas-holdup RMSE ranged from 0.0064 to 0.0068. The changes in the norm ratios indicate that the learned correction responds to the assumed slip parameters, while the evaluated prediction errors remain within the reported ranges.

The corrective term is interpreted as a data-informed adjustment to the prescribed slip relation. Its contribution cannot be assigned uniquely to turbulence, interphase forces, or other mechanisms omitted from the reduced formulation. The present field-norm analysis does not quantify the spatial distribution of the correction, and a separately trained model without Δu_g was not evaluated. Consequently, these results establish its relative global magnitude and sensitivity within the tested parameter ranges, but do not isolate its effect on predictive accuracy.

3.2.2. Gas Velocity Prediction

As shown in Table 4, gas velocity exhibits the highest predictive errors and lowest R^2 values among all evaluated quantities across the surrogate models. This increased difficulty arises because gas velocity, particularly the axial component, is reconstructed from the liquid velocity field, the gas holdup distribution, and an empirical slip-velocity model. Errors in any of these interacting components propagate into the gas velocity prediction. In addition, gas holdup and slip velocity exhibit sharp localized gradients near impeller blades and rising gas plumes, making the gas phase more challenging to resolve than the comparatively smoother liquid velocity field. The purely supervised NN exhibits the weakest performance for gas velocity prediction. At $N_c = 400$ and $N_c = 500$ rpm, R^2 values range from approximately 0.67 to 0.77, with normalized L_1 errors between 58–75%. RMSE values (0.26 – 0.28 m/s) further indicate significant deviation from the CFD reference, particularly near impeller discharge zones and along the shaft. Incorporating boundary-condition enforcement through the BINN yields substantial improvement over the purely supervised NN. RMSE is reduced by approximately 46–56%, normalized errors decrease to about 35%, and R^2 increases to 0.83–0.86, indicating improved near-wall and near-impeller behavior. Nevertheless, because the governing conservation laws are not embedded in the formulation, the BINN continues to underpredict gas circulation strength in the bulk region. The PINN delivers the most accurate gas velocity predictions, with RMSE reduced to 0.08–0.11 m/s, R^2 increasing to approximately 0.95, and normalized L_1 and L_2 errors falling below 9%. By incorporating a gas-transport residual and a reduced liquid momentum constraint during training, the PINN produces smoother and more physically consistent gas velocity fields, particularly in regions with strong velocity gradients near the impellers and gas plumes. Because the reference CFD simulations were verified through a mesh-independence study and previously validated against experimental measurements of gas holdup and gassed power consumption, the observed improvements are primarily attributed to the incorporation of physical constraints into the learning process rather than to the suppression of numerical artifacts in the CFD reference data.

The larger gas velocity error relative to pressure may be partly explained by the reduced gas-phase representation. Gas velocity is reconstructed from the predicted liquid velocity, a gas-holdup-dependent slip relation, and a learned corrective velocity; errors in these terms can therefore propagate into the reconstructed field. Such effects may be more pronounced near the impellers, where gas holdup, shear, and gas–liquid interactions vary rapidly. The gas-continuity residual constrains the divergence of the predicted gas

flux, but the PINN does not solve a separate gas-phase momentum equation. Pressure, by comparison, is directly constrained through its gradient in the reduced momentum residual and through the pressure-gauge and anchor treatment. These differences provide a plausible explanation for the lower pressure errors and the concentration of gas velocity errors near the impellers, although the present analysis does not isolate the contribution of each mechanism.

3.2.3. Liquid Velocity Prediction

As presented in Table 4, the liquid velocity field was predicted more accurately than the gas velocity field by all three surrogate models at both $N_c = 400$ and 500 rpm. The supervised NN achieved coefficients of determination of approximately 0.84 with normalized L_1 errors of 54–56%. The smoother and more coherent structure of the liquid flow allows the network to capture the dominant circulation patterns using supervised data alone. Incorporating boundary-condition constraints in the BINN further improved the predictions, increasing R^2 to approximately 0.91 and reducing the normalized L_1 error to 24–26%, while slightly lowering the RMSE. The PINN consistently provided the highest accuracy, reducing the RMSE to 0.064–0.069 m/s, increasing R^2 to 0.96–0.97, and reducing the normalized L_1 error to below 10% at both operating conditions. These results demonstrate that incorporating momentum conservation into the learning process significantly improves liquid velocity predictions, even for relatively smooth flow fields.

3.2.4. Pressure Field Prediction

The results presented in Table 4 show that pressure prediction exhibits the greatest benefit from incorporating physics-based constraints. At both $N_c = 400$ and 500 rpm, the supervised NN and BINN produced similar RMSE values (208–213 Pa), with R^2 ranging from 0.86 to 0.94. Although the BINN improved R^2 relative to the NN, enforcing boundary conditions alone was insufficient to accurately constrain the pressure field, which is primarily governed by the global momentum balance and incompressibility constraint. In contrast, the PINN substantially improved pressure prediction, reducing the RMSE to 72–88 Pa, increasing R^2 to 0.98–0.99, and reducing the normalized L_1 error to 3–4%. These results demonstrate that incorporating the reduced momentum and liquid-divergence constraints improved pressure prediction in the evaluated cases.

3.2.5. Gas-Holdup Prediction

Table 4 shows that the PINN predicted the CFD gas-holdup field more accurately than the NN and BINN at both evaluated speeds. At $N_c = 400$ rpm, the gas-holdup RMSE decreased from 0.0153 for the supervised NN to 0.0091 for the BINN and 0.0042 for the PINN. The corresponding R^2 values were 0.8118, 0.8912, and 0.9329, while the normalized L_1 errors were 57.90%, 28.12%, and 9.18%, respectively. The PINN's normalized L_2 error was 8.67%. At $N_c = 500$ rpm, the RMSE values for the NN, BINN, and PINN were 0.0214, 0.0115, and 0.0044, respectively, and their R^2 values were 0.8048, 0.8712, and 0.9222. The normalized L_1 error decreased from 56.26% for the NN and 26.03% for the BINN to 8.99% for the PINN; the PINN's normalized L_2 error was 8.50%. The comparison shows better agreement between the PINN and the CFD gas-holdup fields at these speeds, consistent with the contribution of its additional reduced physics constraints. These field-error metrics do not establish that the PINN independently conserves the prescribed global gas throughput.

3.2.6. Sensitivity of the Corrective Gas Velocity Component

The sensitivity of the learned corrective gas velocity, Δu_g , was examined by independently varying its regularization weight, $\lambda_{\Delta u}$, the terminal rise velocity, u_t , and the hindrance exponent, n_h , by $\pm 20\%$ relative to their baseline values. Table 5 reports the resulting field-norm ratios of the correction to the prescribed slip velocity, R_{slip} , and of the liquid velocity, R_l . In the baseline case, $R_{slip} = 0.1460$ and $R_l = 0.0438$. Decreasing $\lambda_{\Delta u}$ by 20% increased these ratios to 0.1668 and 0.0490, whereas increasing it by 20% reduced them to 0.1251 and 0.0386. The direction of this change is consistent with the regularization term's role in limiting the learned correction. The correction also responded to changes in the prescribed slip relation. Across the u_t and n_h perturbations, R_{slip} ranged from 0.1147 to 0.1877, and R_l ranged from 0.0376 to 0.0522. Thus, over the evaluated field, the norm of the learned correction remained below 19% of the prescribed slip-field norm and below 6% of the liquid velocity field norm for all cases examined. These are global field-norm ratios and do not establish that the correction is small at every location. They also do not establish how the parameter perturbations affect the gas velocity prediction errors reported in Table 4. The results therefore characterize the relative magnitude and sensitivity of the learned correction, rather than its independent contribution to predictive accuracy.

Table 5. Sensitivity of the learned corrective gas velocity and predictive errors at $N_c = 300$ rpm to $\pm 20\%$ perturbations in the corrective-velocity regularization weight, terminal rise velocity, and hindrance exponent.

Parameter Case	R_{slip}	R_l	Gas Velocity RMSE (m/s)	α_g RMSE	Change in Gas Velocity RMSE (%)
Baseline	0.1460	0.0438	0.0379	0.0065	-
$\lambda_{\Delta u} - 20\%$	0.1668	0.0490	0.0369	0.0064	-2.64
$\lambda_{\Delta u} + 20\%$	0.1251	0.0386	0.0400	0.0067	5.54
$u_t - 20\%$	0.1877	0.0522	0.0421	0.0068	11.08
$u_t + 20\%$	0.1147	0.0376	0.0411	0.0066	8.44
$n_h - 20\%$	0.1668	0.0480	0.0400	0.0067	5.54
$n_h + 20\%$	0.1356	0.0407	0.0411	0.0068	8.44

$R_{slip} = \|\Delta u_g\|_2 / \|u_{slip}\|_2$ and $R_l = \|\Delta u_g\|_2 / \|u_l\|_2$, where the norms are evaluated over the considered spatial field. Δu_g is the learned corrective gas velocity component, u_{slip} is the prescribed hindered-rise slip velocity, and u_l is the liquid velocity. The L_2 norms are evaluated over the considered spatial field. $\lambda_{\Delta u}$ denotes the regularization weight applied to the corrective gas velocity component, u_t is the terminal rise velocity used in the prescribed slip relation, n_h is the hindrance exponent, and α_g is the gas volume fraction (gas holdup).

3.3. Qualitative Flow-Field Comparison

To complement the quantitative error metrics, a qualitative comparison was conducted to assess how well the predictive models capture the flow pattern and physical features of the velocity and pressure fields. Representative contours were compared against CFD reference solutions at $N_c = 400$ rpm and $N_c = 500$ rpm.

3.3.1. Axial Gas Velocity Contours

Figure 4 compares the contours of the axial gas velocity component, $u_{g,y}$, on the vertical mid-plane ($z = 0$) predicted by the CFD model and the three surrogate approaches.

At $N_c = 400$ rpm, the CFD solution exhibits two distinct gas-rich upflow regions originating from the discharge zones of the lower and upper impellers. These structures result from the combined effects of impeller-induced axial pumping and buoyancy-driven gas rise. The lower impeller initiates gas dispersion and axial transport, whereas the upper impeller redistributes the gas upward, producing the characteristic inter-impeller circulation pattern of aerated double-impeller stirred tanks.

The supervised NN reproduces the overall symmetry and approximate magnitude of the axial gas velocity but excessively smooths the impeller discharge jets and gas-rich upflow regions, resulting in blurred inter-impeller structures and weakened velocity gradients. These deficiencies indicate that a purely data-driven model cannot adequately capture the nonlinear coupling among gas transport, buoyancy, and impeller-induced momentum transfer. Incorporating boundary conditions through the BINN improves the near-wall predictions and partially restores flow symmetry; however, axial transport and circulation in the central region remain underestimated, indicating that boundary constraints alone are insufficient to recover the internal gas dynamics governed by the conservation equations. In contrast, the PINN closely reproduces the CFD solution, accurately capturing the location and intensity of the gas-rich upflow regions, preserving the sharp velocity gradients near the impeller discharge zones, and resolving the inter-impeller circulation without introducing artificial diffusion or numerical oscillations.

At $N_c = 500$ rpm, the overall flow topology is preserved, but the gas-rich upflow regions become stronger and more vertically elongated owing to the higher impeller speed. Under these more demanding conditions, the limitations of the NN and BINN become increasingly evident. Both models underestimate the strength of the gas-rich upflow regions and excessively smooth the axial velocity gradients, particularly within the inter-impeller region. Although the BINN continues to improve the near-wall predictions relative to the NN, it remains unable to reproduce the strengthened internal circulation. In contrast, the PINN maintains excellent agreement with the CFD solution by accurately resolving the intensified axial transport and enhanced inter-impeller coupling. The agreement at both operating conditions indicates that incorporating reduced physics constraints improved the PINN's prediction of the gas–liquid flow structures observed in the CFD reference fields.




Figure 4. Axial gas velocity contours for various surrogate models at $N_c = 400$ and $N_c = 500$ rpm.

3.3.2. Radial Gas Velocity at the Lower Impeller Plane

Figure 5 compares contours of the radial gas velocity on the horizontal x – z plane passing through the lower impeller ($y = 0.18$ m) at $N_c = 400$ and 500 rpm. The CFD results exhibit a distinct four-lobed pattern aligned with the impeller blades, corresponding to regions of strong outward gas transport generated by blade-induced pumping. Peak radial velocities occur within the blade-swept zones and gradually decrease toward the tank

wall as the gas is redistributed into the bulk flow. The purely supervised NN reproduces the overall velocity magnitude but fails to resolve the blade-induced flow structure. The lobes become broadened and partially merged, peak velocities are underpredicted, and the sharp gradients near the impeller tips are excessively smoothed, leading to a loss of jet confinement. These deficiencies demonstrate the difficulty of learning highly localized, geometry-dependent multiphase flow features from data alone. The BINN improves the overall symmetry and reduces near-wall distortions by enforcing boundary conditions at the impeller, shaft, and tank walls. However, it continues to underpredict the intensity and spatial confinement of the radial jets, indicating that boundary constraints alone are insufficient to accurately capture the internal gas–impeller interactions governing radial gas transport. The PINN closely reproduces the CFD solution at both operating conditions. It accurately captures the orientation, intensity, and spatial confinement of the four-lobed pattern while preserving the sharp velocity gradients near the impeller tips. The transition from the impeller discharge region to the surrounding bulk flow is also well represented. These results demonstrate that incorporating momentum conservation and gas-continuity constraints enables more accurate reconstruction of impeller-driven gas transport than purely data-driven or boundary-informed models.




Figure 5. Radial contours of gas velocity at $y = 0.18$ m (bottom impeller) at $N_c = 400$ rpm and $N_c = 500$ rpm.

3.3.3. Effect of Increasing Impeller Speed on Velocity Field

Figures 4 and 5 show that increasing the impeller speed from $N_c = 400$ to 500 rpm enhances gas transport in both the axial and radial directions. The CFD results exhibit stronger impeller discharge jets, steeper velocity gradients near the blade tips, and increased interaction between the upper and lower impellers, reflecting greater momentum transfer and gas–liquid coupling at the higher rotational speed. The supervised NN and BINN preserve the overall flow topology but exhibit larger deviations as impeller-driven transport becomes stronger. Both models increasingly underpredict the strength of the axial and radial jets and oversmooth high-gradient regions, particularly in the inter-impeller region, indicating reduced predictive accuracy under more demanding flow conditions. In contrast, the PINN maintains high accuracy at both operating conditions. The intensified axial upflow regions, blade-induced radial jets, and their spatial localization are accurately reproduced, together with the associated velocity gradients. These results show that incorporating reduced momentum and gas-transport constraints helps the PINN reproduce the CFD flow structures as impeller speed increases.

3.3.4. Pressure Field Contours

Figures 6 and 7 compare the axial and radial pressure contours predicted by the supervised NN, BINN, and PINN with the CFD reference solutions at $N_c = 400$ and 500 rpm. The axial plane illustrates the pressure distribution along the vessel height, while the radial plane ($y = 0.18$ m) intersects the lower impeller region, where blade-induced pressure gradients are most pronounced. The CFD results in Figure 6 reveal well-defined low-pressure regions around both impellers at both operating conditions. These pressure minima arise from blade-induced suction and the associated fluid acceleration near the impeller discharge zones. The axial contours also exhibit a gradual decrease in pressure from the tank bottom to the upper region, reflecting the combined effects of hydrostatic pressure and impeller-driven flow. The supervised NN captures the overall pressure distribution but fails to resolve the localized extrema near the impellers. Pressure gradients within the blade-swept regions are substantially smoothed in both the axial and radial planes, particularly around the lower impeller. This limitation becomes more pronounced at $N_c = 500$ rpm, where the pressure minima are increasingly attenuated and the blade-induced pressure field becomes less distinct. These deficiencies arise because the supervised NN relies solely on data fitting and does not enforce the momentum conservation and incompressibility constraints that govern pressure–velocity coupling in incompressible gas–liquid flows. The BINN provides moderate improvement by more accurately reproducing the large-scale pressure distribution and the impeller-induced low-pressure regions. Enforcing boundary conditions improves predictions near the walls and impeller surfaces while enhancing the overall symmetry of the solution. However, attenuation of the pressure extrema and distortion of the blade-induced pressure field remain evident, indicating that boundary constraints alone are insufficient to accurately reconstruct the internal pressure distribution governed by global momentum balance. Both the NN and BINN also exhibit noticeable errors along the shaft centerline, where geometric confinement and localized pressure variations are only weakly constrained by the boundary conditions and sparsely represented in the training data. The PINN achieves the closest agreement with the CFD reference at both operating conditions. The axial and radial contours accurately reproduce the location, extent, and magnitude of the impeller-induced low-pressure regions while preserving the global axial pressure gradient and flow symmetry. Predictions along the shaft centerline are also significantly improved. These results indicate that incorporating reduced momentum and liquid-divergence constraints improves agreement with the CFD pressure field, including in regions with strong pressure gradients and impeller-driven flow.




Figure 6. Pressure contours on the vertical mid-plane ($z = 0$) obtained from CFD and surrogate models (supervised NN, BINN, and PINN) at impeller rotational speeds $N_c = 400$ rpm and $N_c = 500$ rpm.




Figure 7. Pressure contours on the radial plane at $y = 0.18$ m (lower impeller plane) obtained from CFD and surrogate models (NN, BINN, and PINN) at $N_c = 400$ rpm and $N_c = 500$ rpm.

3.3.5. Effect of Increasing Impeller Speed on Pressure Field

A comparison between $N_c = 400$ and 500 rpm further illustrates the sensitivity of the pressure field to increasing impeller speed. As the rotational speed increases, the CFD solution exhibits steeper pressure gradients and more pronounced low-pressure regions around the impellers, particularly near the lower impeller where circulation and gas entrainment are strongest. These features reflect enhanced momentum transfer and centrifugal effects at higher rotational speeds. The differences between the surrogate models become more evident under these conditions. The supervised NN increasingly underpredicts the pressure minima and oversmooths the localized gradients near the impeller blades, indicating reduced accuracy as the flow becomes more strongly driven. The BINN responds more effectively to the increased impeller speed but continues to distort the localized pressure distribution, particularly on the radial plane. In contrast, the PINN maintains close agreement with the CFD reference at both operating conditions. The location, extent, and magnitude of the impeller-induced low-pressure regions are accurately reproduced, together with the associated pressure gradients. These results demonstrate that enforcing momentum conservation and incompressibility enables the PINN to reliably predict pressure fields under increasingly demanding operating conditions.

3.4. Axial Gas Velocity Profiles Along the Tank Height

Figure 8 presents axial gas velocity profiles ($u_{g,y}$) along the vessel height for $N_c = 400$ and 500 rpm, comparing the CFD and PINN predictions. The profiles are extracted along a vertical line at m , which passes through the impeller-swept regions and therefore captures the dominant axial gas transport. At both operating conditions, the CFD profiles exhibit two distinct velocity peaks near the lower (m) and upper (m) impellers, corresponding to the upward gas transport generated by impeller pumping and buoyancy. Between the impellers, the axial velocity decreases, indicating gas redistribution within the inter-impeller interaction zone, a characteristic feature of double-impeller stirred tanks.

Increasing the impeller speed from $N_c = 400$ to 500 rpm strengthens the axial gas transport, resulting in higher peak velocities and steeper gradients near both impellers, with the largest increase occurring around the upper impeller. The PINN accurately reproduces these trends, capturing the peak locations, magnitudes, and overall profile shape at both operating conditions. Minor smoothing near the velocity extrema and small deviations close to the tank bottom and free surface are observed, where the gas velocity is low and relative errors become more sensitive. Despite the absence of dedicated gas-phase boundary-loss terms, the predicted profiles do not exhibit abrupt boundary-localized variations indicative of obvious spurious gas velocity boundary layers within the evaluated cases. However, because the present formulation does not independently enforce the

inlet-to-outlet gas flux, strict global gas-phase flux conservation cannot be inferred from these profile comparisons alone. The close agreement between the PINN and CFD profiles demonstrates the model's ability to accurately predict the axial gas velocity distribution and its dependence on impeller speed. These results further confirm that incorporating momentum conservation and gas continuity into the training process improves the accuracy and physical consistency of gas-phase predictions.




Figure 8. Axial gas velocity profiles at $x = 0.09$ m for $N_c = 400$ rpm and $N_c = 500$ rpm.

3.5. Prediction at Unseen Intermediate Operating Conditions

3.5.1. Gas Velocity Field Interpolation

To evaluate the interpolation capability of the PINN between the discrete training conditions, predictions were performed at intermediate impeller speeds of 350 and 450 rpm, which lie within the training range of 200–500 rpm but were excluded from the training dataset. These conditions therefore assess the model's ability to interpolate continuously with respect to impeller speed. Because the NN and BINN were not evaluated at these intermediate conditions, the present analysis assesses only the interpolation capability of the PINN and should not be interpreted as demonstrating superior interpolation relative to the baseline models. Figure 9 compares the CFD and PINN contours of the axial gas velocity on the vertical mid-plane ($z = 0$) and the lower-impeller radial plane ($y = 0.18$ m). At both interpolation speeds, the PINN closely reproduces the CFD flow patterns in both planes. In the axial plane, the dual-impeller upflow regions, inter-impeller interaction zone, and vertical gas transport are accurately captured. The increase in axial gas transport from 350 to 450 rpm is also well reproduced, with stronger axial jets and steeper velocity gradients near the impeller discharge regions. In the radial plane, the characteristic four-lobed pattern aligned with the impeller blades is preserved, together with the magnitude and spatial confinement of the blade-induced radial jets and the gradual redistribution of gas toward the bulk flow. The close agreement across both planes demonstrates that the PINN learns a smooth and physically consistent relationship between the flow field and impeller speed within the investigated operating range. Rather than reproducing only the discrete training cases, the model accurately captures the continuous variation in gas transport with impeller speed, enabling reliable interpolation at unseen intermediate operating conditions.

Quantitatively, the PINN achieves normalized L_1 and L_2 errors of 13.7% and 13.9%, respectively, for the axial gas velocity at $N_c = 350$ and 450 rpm, the corresponding errors

decrease to 11.5% (L_1) and 11.0% (L_2), indicating improved agreement with the CFD reference. Although these errors are higher than those obtained at the trained operating conditions (Table 4), the predicted gas velocity fields remain physically consistent. The PINN accurately preserves the coherent axial gas plumes above both impellers and the smooth transition between the impeller-dominated regions and the surrounding bulk flow. The remaining discrepancies are primarily confined to high-gradient regions near the impeller blades, where strong gas–liquid coupling and slip-induced effects amplify small deviations in the reconstructed gas velocity.

The present study focuses on predicting the hydrodynamic fields rather than derived engineering quantities. Consequently, the impact of the interpolation errors on quantities such as the gas–liquid interfacial area and volumetric mass-transfer coefficient ($k_L a$) was not evaluated. These quantities depend on additional models describing bubble size distribution, interfacial area, and mass transfer, which are beyond the scope of the present reduced-order framework. Quantifying the propagation of surrogate-model errors to mass-transfer predictions will be the subject of future work.

The increase in gas velocity error to approximately 11–14% at the unseen intermediate speeds indicates a limitation in interpolation accuracy, particularly near the impellers. The largest deviations are observed in the blade-swept and impeller-discharge regions, where gas transport varies rapidly with rotational speed. The reduced slip representation may contribute to these errors, although its individual effect was not isolated. Future training datasets could be enriched at intermediate speeds and in regions with larger validation errors. Targeted sampling near the impellers and residual-driven collocation could also be evaluated as ways to improve local accuracy.




Figure 9. Comparison of CFD and PINN gas velocity contours at two unseen impeller speeds ($N_c = 350$ rpm and $N_c = 450$ rpm) on axial and radial planes.

3.5.2. Pressure Field Interpolation

Figure 10 compares the CFD and PINN pressure contours on the vertical mid-plane ($z = 0$) and the lower-impeller radial plane ($y = 0.18$ m) at the interpolation speeds of $N_c = 350$ and 450 rpm. On the axial plane, the CFD solution exhibits the expected vertical pressure stratification, with higher pressure near the tank bottom and lower pressure toward the upper region, together with distinct low-pressure regions around both impellers resulting from blade-induced suction. As the impeller speed increases from 350 to 450 rpm, these low-pressure regions deepen and expand, reflecting stronger blade-induced acceleration and increased inertial effects.

The PINN closely reproduces these features at both operating conditions. The location, extent, and relative magnitude of the low-pressure regions are accurately captured while preserving the global axial pressure gradient. Minor smoothing is observed in the immediate vicinity of the blade-swept regions; however, the overall pressure distribution remains in close agreement with the CFD reference.

On the radial plane ($y = 0.18$ m), the CFD contours reveal a four-lobed pressure pattern aligned with the impeller blades. These localized pressure minima become more pronounced at 450 rpm, consistent with the increased inertial effects associated with higher rotational speed. The PINN accurately reconstructs the four-lobed structure and its evolution with impeller speed while preserving the symmetry and spatial localization of the blade-induced pressure field. Slight attenuation of the peak pressure minima is observed, particularly at 450 rpm, reflecting the inherent difficulty of reconstructing sharp localized gradients from finite training data. Nevertheless, the dominant blade-resolved pressure structures are retained.

These results demonstrate that the PINN provides accurate and physically consistent pressure predictions at unseen intermediate operating conditions. The close agreement with the CFD reference indicates that the reduced momentum and liquid-divergence constraints help the PINN interpolate the pressure field at the evaluated intermediate speeds.




Figure 10. Comparison of CFD and PINN pressure contours at two unseen impeller speeds ($N_c = 350$ rpm and $N_c = 450$ rpm) on axial and radial planes.

3.5.3. Gas-Holdup Interpolation

Figure 11 compares the local gas-holdup distributions obtained from CFD and the PINN at the intermediate impeller speeds of $N_c = 350$ and 450 rpm. The axial and radial contours show close qualitative agreement between the CFD and PINN predictions at both operating conditions, indicating that the model reproduces the principal spatial features of the gas distribution.

At $N_c = 350$ rpm, elevated gas holdup is concentrated primarily in the impeller regions and along the associated discharge paths, while lower values occur throughout much of the surrounding bulk liquid. The PINN reproduces the location and spatial extent of these higher-holdup regions, including the localized structures surrounding the impellers. Similar agreement is observed in the horizontal plane, where the nonuniform distribution around the impeller and the regions of locally elevated gas holdup are captured by the PINN. Minor differences are visible in some localized regions, particularly near the impeller and within areas characterized by sharp spatial variations in gas holdup.




Figure 11. Comparison of local gas-holdup contours obtained from CFD and the PINN at $N_c = 350$ and 450 rpm.

Increasing the impeller speed to $N_c = 450$ rpm modifies both the magnitude and spatial distribution of the local gas holdup. The regions of elevated gas concentration around the impellers become more pronounced, particularly in the axial plane, while the horizontal contours exhibit a clearer distribution of high-holdup regions around the impeller. The PINN closely reproduces these changes, including the location of the major high-holdup zones and the surrounding lower-holdup regions. This agreement indicates that the model captures the variation in the gas-holdup field between the discrete rotational speeds used for training rather than merely reproducing the distributions at the trained operating conditions.

The prediction of gas holdup is particularly important in the present framework because α_g contributes directly to the effective mixture properties used in the reduced momentum residual and to the reconstruction of the gas-phase velocity. Therefore, discrepancies in the predicted gas-holdup field can propagate to the reconstructed gas velocity and partly contribute to the larger local errors observed for the gas-phase velocity components. The close correspondence between the CFD and PINN contours in Figure 11 provides qualitative support for the ability of the surrogate to represent the spatial gas distribution at intermediate operating conditions. Quantitative gas-holdup errors and volume-averaged values should be reported alongside these contours to further assess the magnitude of the remaining discrepancies.

The qualitative agreement observed in Figure 11 is further supported by the quantitative gas-holdup results summarized in Table 6. At $N_c = 350$ rpm, the CFD and PINN volume-averaged gas holdups were 0.0310 and 0.0295, respectively, corresponding to a relative deviation of approximately 4.8%. At $N_c = 450$ rpm, the corresponding values were 0.0430 and 0.0412, with a relative deviation of approximately 4.2%. The normalized L_1 errors of the local gas-holdup field were 9.2% and 7.8% at $N_c = 350$ and 450 rpm, respectively, while the corresponding RMSE values remained below 0.005. These results indicate that the PINN reproduces not only the principal spatial distribution of gas holdup observed in the CFD contours but also its overall magnitude at the intermediate operating conditions.

The local-field errors are higher than the deviations in volume-averaged gas holdup, indicating that localized discrepancies partially compensate when integrated over the vessel volume. This behavior is consistent with the contours in Figure 11, where the largest differences between the CFD and PINN predictions occur primarily in localized regions near the impellers and their discharge zones, while the broader gas-distribution pattern remains similar.

Since gas holdup enters the effective properties and gas velocity reconstruction in the reduced PINN formulation, these localized discrepancies may also contribute to the comparatively larger errors observed in the reconstructed gas-phase velocity field.

Table 6. Quantitative comparison of gas-holdup predictions obtained from CFD and the PINN at the intermediate impeller speeds of $N_c = 350$ and 450 rpm.

N_c (rpm)	CFD α_g	PINN α_g	Relative Error (%)	Normalized L_1 Error (%)	RMSE
350	0.093	0.098	5.38	9.2	0.0051
450	0.121	0.128	5.79	7.8	0.0072

3.5.4. Quantitative Interpolation Errors

While the contour comparisons demonstrate close agreement between the CFD and PINN pressure fields, quantitative error metrics provide a more rigorous assessment of the interpolation performance. The pressure field is interpolated more accurately than the gas velocity field. At $N_c = 350$ rpm, the PINN achieves normalized L_1 and L_2 pressure errors of 4.48% and 6.74%, respectively. At $N_c = 450$ rpm, the corresponding errors remain low at 4.97% and 9.22%, respectively. Thus, the reported pressure errors span 4.48–9.22%, depending on speed and error metric. The pressure contours also show that the PINN reproduces the main distribution and impeller-associated low-pressure regions. The reduced momentum residual contains a pressure-gradient term, but these results do not isolate the effect of that constraint from the supervised CFD data or other loss terms.

Compared with the pressure field, gas velocity exhibits larger interpolation errors, particularly in the high-shear regions near the impeller blades, where slip dynamics and complex gas–liquid interactions are associated with larger local deviations. Nevertheless, the PINN consistently preserves the dominant flow structures and the dependence of both axial and radial gas transport on impeller speed. Although interpolation errors increase locally in the blade-swept regions, the predicted flow fields remain physically consistent and accurately reproduce the principal hydrodynamic features throughout the investigated operating range. This localization of the interpolation errors is consistent with the limitations of the reduced physics formulation discussed in Section 3.2.1, particularly in regions characterized by strong shear and complex gas–liquid momentum coupling. These results demonstrate that incorporating reduced conservation-based physics constraints into the training process effectively constrains the solution space and enables reliable interpolation between the discrete training conditions.

3.6. Computational Efficiency Evaluation

The computational cost of the proposed framework comprises generation of the reference CFD database, PINN training, and subsequent inference. Each three-dimensional reference CFD simulation required approximately 100–140 h of wall-clock time using 32 CPU cores. The four simulations at $N_c = 200, 300, 400$, and 500 rpm therefore required approximately 400–560 h in aggregate when their run times are summed. PINN training required a further 2–3 h on a workstation equipped with dual Intel® Xeon® Gold 6148 CPUs (2.40 GHz), 256 GB RAM, and an NVIDIA GTX 1060 (6 GB). On the same workstation, inference at 200,000 spatial locations required approximately 2–3 h for one operating condition. The computational benefit depends on whether the reference CFD database is already available. If it is, training adds approximately 2–3 h, after which each additional operating condition requires approximately 2–3 h of PINN inference rather than 100–140 h for a new CFD simulation. If the four CFD cases are generated specifically for surrogate development, their computational cost must also be counted before assessing the overall

benefit. Using the midpoints of the reported ranges, the four reference CFD simulations required a combined 480 h, PINN training required 2.5 h, and inference required 2.5 h per operating condition. By comparison, an independent CFD simulation required 120 h per condition. When the four reference conditions are counted among the total conditions evaluated, the surrogate's aggregate computational cost becomes lower at the fifth condition: four reference conditions plus one additional condition. This comparison uses the sum of CFD run times; the elapsed time to produce the database would differ if the cases were run concurrently. The benefit increases as the trained model is used for further evaluations within the investigated operating range.

This reduction in evaluation time is accompanied by limits in local accuracy. Gas velocity errors are larger than pressure errors, particularly near the impellers, where phase-slip gradients and gas–liquid interactions are complex. The PINN approximates the CFD-derived fields using the reduced physics constraints described in Section 2.4; it does not reproduce the complete Euler–Euler/RANS formulation. The model nevertheless captures the dominant flow features and their variation with impeller speed. Its use in parametric studies, optimization, real-time monitoring, or digital-twin applications remains to be demonstrated.

The PINN is intended as a complementary surrogate rather than a replacement for CFD or experimental validation. Its demonstrated use is interpolation within the investigated operating domain. Predictions outside the 200–500 rpm training range were not validated and should be checked against CFD and, where possible, experimental measurements. Because gas flow rate and fluid properties were held constant and were not network inputs, the results also do not establish accuracy when these quantities change. Extending the surrogate to such conditions would require reference data spanning those variables and their inclusion as model inputs.

A further limitation is that the framework represents only the steady hydrodynamic states described by the multiple-reference-frame (MRF) CFD simulations. It has not been evaluated for startup, shutdown, abrupt changes in rotational speed, or periodic blade–baffle interactions. Studying unsteady operation would require transient reference data, potentially from sliding-mesh CFD, and a model that includes time as an input. A time-dependent PINN could incorporate transient physics constraints, while long short-term memory (LSTM) or gated recurrent unit (GRU) networks could be investigated for representing temporal flow dynamics.

4. Conclusions

This study developed a PINN surrogate for steady three-dimensional gas–liquid hydrodynamics in an aerated double-impeller stirred tank. A supervised neural network and a boundary-informed neural network with the same architecture were evaluated to assess the effects of adding boundary and reduced physics constraints. At $N_c = 400$ and 500 rpm, the PINN outperformed the NN and BINN, with normalized liquid- and gas velocity errors below 10% and pressure errors of approximately 3–4%. It also better preserved impeller-induced gradients, particularly in high-shear regions. At the unseen intermediate speeds of 350 and 450 rpm, the PINN reproduced the principal flow structures, although gas velocity interpolation errors of approximately 11–14% exceeded those for pressure. The four reference CFD cases required approximately 400–560 h in total when their run times were summed, followed by 2–3 h of PINN training. Once trained, the PINN required approximately 2–3 h to evaluate 200,000 spatial locations for one operating condition, compared with 100–140 h for a CFD simulation. The benefit of repeated surrogate evaluations must therefore be assessed alongside the initial cost of generating the CFD training data. The PINN is a complementary surrogate for the investigated geometry, fluid system, gas

flow rate, and rotational-speed range. Its reduced physics constraints do not reproduce the complete CFD formulation. Predictions for new geometries, operating regimes, or transient conditions require further CFD evaluation and, where possible, experimental validation. Future work could extend the training domain and incorporate transient hydrodynamics and coupled mass-transfer models.

Author Contributions: Conceptualization, F.S., A.L. and F.E.-M.; methodology, F.S., A.L. and F.E.-M.; software, F.S.; validation, F.S.; formal analysis, F.S.; investigation, F.S.; data curation, F.S.; writing—original draft preparation, F.S.; writing—review and editing, A.L. and F.E.-M.; visualization, F.S.; supervision, A.L. and F.E.-M.; project administration, F.E.-M. and A.L.; funding, A.L. and F.E.-M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) through the Discovery Grant Program, grant numbers RGPIN-2019-04644 and RGPIN-2019-05644.

Data Availability Statement: Dataset available on request from the authors.

Acknowledgments: The authors gratefully acknowledge Ansys for providing the necessary licenses used in this research. Computations were performed on the Trillium cluster at the SciNet HPC Consortium. SciNet is funded by Innovation, Science and Economic Development Canada; the Digital Research Alliance of Canada; the Ontario Research Fund: Research Excellence; and the University of Toronto.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

Latin Symbols

g	Gravitational acceleration	m s^{-2}
$h^{(l)}$	Output of layer l	—
L	Total loss function	—
L_1	Normalized mean absolute error	—
L_2	Normalized root mean square error	—
L_{BC}	Boundary-condition loss	—
L_{data}	Supervised data loss	—
L_{PDE}	Governing equation residual loss	—
n_h	Hindrance exponent	—
N_c	Impeller rotational speed	rpm
N_c^*	Nondimensional impeller speed $N_c/N_{c,\text{ref}}$	—
$N_\theta(\cdot)$	Neural network mapping	—
p	Pressure	Pa
p^*	Nondimensional pressure	—
r	Radial position vector	m
R	Residual vector of governing equations	—
RMSE	Root mean square error	field-dependent
u_g	Gas-phase velocity vector	m s^{-1}
u_l	Liquid-phase velocity vector	m s^{-1}
$u_{l,x}, u_{l,y}, u_{l,z}$	Liquid velocity components	m s^{-1}
u_{slip}	Prescribed slip velocity	m s^{-1}
u_t	Terminal bubble rise velocity	m s^{-1}
$W^{(l)}$	Weight matrix of layer l	—
x, y, z	Spatial coordinates	M
X_{BC}	Boundary collocation points	—
X_{data}	Labeled data points	—
X_{PDE}	Interior collocation points	—

Greek Symbols

α_g	Gas holdup	–
α_{raw}	Unconstrained network holdup output	–
β_1, β_2	Adam momentum parameters	–
Δu_g	Learned corrective gas velocity component	m s^{-1}
ϵ	Numerical stability parameter	–
η	Learning rate	–
λ_{BC}	Boundary-loss weight	–
λ_{data}	Data loss weight	–
λ_{PDE}	PDE residual loss weight	–
μ_g	Gas viscosity	Pa s
μ_l	Liquid viscosity	Pa s
μ_{eff}	Effective mixture viscosity	Pa s
ρ_g	Gas density	kg m^{-3}
ρ_l	Liquid density	kg m^{-3}
ρ_{eff}	Effective mixture density	kg m^{-3}
$\sigma(\cdot)$	Activation function	–
θ	Trainable neural network parameters	–
ω	Angular velocity	Rad s^{-1}
ϕ	Generic scalar field	–
ϕ_{ref}	Reference scaling quantity	–

References

- Joshi, J.B.; Nere, N.K.; Rane, C.V.; Murthy, B.N.; Mathpati, C.S.; Patwardhan, A.W.; Ranade, V.V. CFD simulation of stirred tanks: Comparison of turbulence models. Part I: Radial flow impellers. *Can. J. Chem. Eng.* **2011**, *89*, 23–82. [\[CrossRef\]](#)
- Spier, M.R.; Lopes da Silva, V.; Bianchi Pedroni Medeiros, M.A.; Rossi, S.C. Application of different types of bioreactors in bioprocesses. In *Bioreactors: Design, Properties and Applications*; Nova Science Publishers: New York, NY, USA, 2011.
- Zamani Abyaneh, E.; Ein-Mozaffari, F.; Lohi, A. Investigating gas dispersion and solid suspension in a coaxial stirred tank containing non-Newtonian fluids using coupled tomography and pressure measurements. *Powder Technol.* **2026**, *467*, 121606. [\[CrossRef\]](#)
- Ranjbari, P.; Ein-Mozaffari, F.; Upreti, S. Analysis of hydrodynamic forces on solid particles in mixing tanks using coupled CFD–DEM method: Influence of impeller pumping direction and speed on mixing dynamics. *Powder Technol.* **2026**, *468*, 121609. [\[CrossRef\]](#)
- Beigzadeh, R.; Soltanian, S.; Azizpour, Y. Enhancing mixing efficiency in multi-impeller tanks through innovative CFD and artificial neural network approaches. *Results Eng.* **2025**, *27*, 106594. [\[CrossRef\]](#)
- Jamshidzadeh, M.; Ein-Mozaffari, F.; Lohi, A. Local and overall gas holdup in an aerated coaxial mixing system containing a non-Newtonian fluid. *AIChE J.* **2020**, *66*, e17016. [\[CrossRef\]](#)
- Sharifi, F.; Behzadfar, E.; Ein-Mozaffari, F. Intensified gas–liquid mixing in bioreactors equipped with a dual coaxial mixer containing biopolymer solutions. *Chem. Eng. Res. Des.* **2023**, *191*, 109–126. [\[CrossRef\]](#)
- Garcia-Ochoa, F.; Gomez, E. Bioreactor scale-up and oxygen transfer rate in microbial processes: An overview. *Biotechnol. Adv.* **2009**, *27*, 153–176. [\[CrossRef\]](#) [\[PubMed\]](#)
- Hinge, S.P.; Kumar, R.; Patwardhan, A.W. Comparison of different impellers for gas dispersion in power-law fluids. *ACS Eng. Au* **2022**, *2*, 75–91. [\[CrossRef\]](#)
- Sharifi, F.; Behzadfar, E.; Ein-Mozaffari, F. Comprehensive investigation of gas hold-up in a double coaxial mixer with shear-thinning fluids exhibiting yield stress: Experimental, numerical, and artificial neural network approaches. *Chem. Eng. Process. Process Intensif.* **2024**, *206*, 110049. [\[CrossRef\]](#)
- Nienow, A.W. Reactor engineering in large-scale animal cell culture. *Cytotechnology* **2006**, *50*, 9–33. [\[CrossRef\]](#) [\[PubMed\]](#)
- Bach, C.; Yang, J.; Larsson, H.; Stocks, S.M.; Gernaey, K.V.; Albæk, M.O.; Krühne, U. Evaluation of mixing and mass transfer in a stirred pilot-scale bioreactor utilizing CFD. *Chem. Eng. Sci.* **2017**, *171*, 19–26. [\[CrossRef\]](#)
- Haringa, C.; Tang, W.; Deshmukh, A.T.; Xia, J.; Reuss, M.; Heijnen, J.J.; Mudde, R.F.; Noorman, H.J. Euler–Lagrange computational fluid dynamics for (bio)reactor scale-down: An analysis of organism lifelines. *Eng. Life Sci.* **2016**, *16*, 652–663. [\[CrossRef\]](#) [\[PubMed\]](#)
- Nienow, A.W.; Nordkvist, M.; Boulton, C.A. Scale-down/scale-up studies leading to improved commercial beer fermentation. *Biotechnol. J.* **2011**, *6*, 911–917. [\[CrossRef\]](#) [\[PubMed\]](#)

15. Ferziger, J.H.; Perić, M. *Computational Methods for Fluid Dynamics*, 3rd ed.; Springer: Berlin/Heidelberg, Germany, 2002.
16. Khopkar, A.R.; Rammohan, A.R.; Ranade, V.V.; Dudukovic, M.P. Gas–liquid flow generated by a Rushton turbine in a stirred vessel: CARPT/CT measurements and CFD simulations. *Chem. Eng. Sci.* **2005**, *60*, 2215–2229. [[CrossRef](#)]
17. Lassila, T.; Manzoni, A.; Quarteroni, A.; Rozza, G. Model order reduction in fluid dynamics: Challenges and perspectives. In *Reduced Order Methods for Modeling and Computational Reduction*; Quarteroni, A., Rozza, G., Eds.; Springer: Cham, Switzerland, 2014; pp. 235–273. [[CrossRef](#)]
18. Mao, R.; Lan, Y.; Liang, L.; Yu, T.; Mu, M.; Leng, W.; Long, Z. Rapid CFD prediction based on machine learning surrogate models in the built environment: A review. *Fluids* **2025**, *10*, 193. [[CrossRef](#)]
19. Barros, P.L.; Ein-Mozaffari, F.; Lohi, A.; Upreti, S. Exploiting the prediction of mass transfer performance in aerated coaxial mixers containing biopolymer solutions using empirical correlations and neural networks. *Can. J. Chem. Eng.* **2025**, *103*, 1258–1275. [[CrossRef](#)]
20. Wang, S.; Teng, Y.; Perdikaris, P. Understanding and mitigating gradient pathologies in physics-informed neural networks. *arXiv* **2020**, arXiv:2001.04536.
21. Brunton, S.L.; Noack, B.R.; Koumoutsakos, P. Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **2024**, *56*, 357–384. [[CrossRef](#)]
22. Meng, X.; Karniadakis, G.E. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *J. Comput. Phys.* **2020**, *401*, 109020. [[CrossRef](#)]
23. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215. [[CrossRef](#)] [[PubMed](#)]
24. Al-Aradi, A.; Correia, A.; Naiff, D.; Jardim, G.; Saporito, Y. Solving nonlinear and high-dimensional partial differential equations via deep learning. *arXiv* **2018**, arXiv:1811.08782.
25. Han, J.; Jentzen, A.; Weinan, E. Solving high-dimensional partial differential equations using deep learning. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 8505–8510. [[CrossRef](#)] [[PubMed](#)]
26. Sarker, I.H. Deep learning: A comprehensive overview on techniques, taxonomy, applications and research directions. *SN Comput. Sci.* **2021**, *2*, 420. [[CrossRef](#)] [[PubMed](#)]
27. Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv* **2021**, arXiv:2010.08895.
28. Lu, L.; Jin, P.; Karniadakis, G.E. DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *Nat. Mach. Intell.* **2021**, *3*, 218–229. [[CrossRef](#)]
29. Trávníčková, V.; Wolff, D.; Dirkes, N.; Elgeti, S.; von Lieres, E.; Behr, M. A model hierarchy for predicting the flow in stirred tanks with physics-informed neural networks. *Adv. Comput. Syst. Eng.* **2024**. [[CrossRef](#)]
30. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [[CrossRef](#)]
31. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [[CrossRef](#)]
32. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics informed deep learning (Part I): Data-driven solutions of nonlinear partial differential equations. *arXiv* **2017**, arXiv:1711.10561.
33. Horiguchi, I.; Shima, K.; Okano, Y. Physics-informed neural networks (PINNs) for high-resolution prediction of shear stress on cells in suspension culture. *AIChE J.* **2025**, *71*, e18853. [[CrossRef](#)]
34. Qiu, R.; Huang, R.; Xiao, Y.; Wang, J.; Zhang, Z.; Yue, J.; Zeng, Z.; Wang, Y. Physics-informed neural networks for phase-field method in two-phase flow. *Phys. Fluids* **2022**, *34*, 052111. [[CrossRef](#)]
35. Zhao, X.; Fan, H.; Lin, G.; Fang, Z.; Yang, W.; Li, M.; Wang, J.; Lu, X.; Li, B.; Wu, K.J.; et al. Multi-objective optimization of a radially stirred tank based on CFD and machine learning. *AIChE J.* **2024**, *70*, e18324. [[CrossRef](#)]
36. Velioglu, M.; Zhai, S.; Rupprecht, S.; Mitsos, A.; Jupke, A.; Dahmen, M. Physics-informed neural networks for dynamic process operations with limited physical knowledge and data. *Comput. Chem. Eng.* **2025**, *192*, 108899. [[CrossRef](#)]
37. Trávníčková, V.; von Lieres, E.; Behr, M. Quantifying data needs in surrogate modeling for flow fields in two-dimensional stirred tanks with physics-informed neural networks. *Phys. Fluids* **2025**, *37*, 117119. [[CrossRef](#)]
38. Jin, X.; Cai, S.; Li, H.; Karniadakis, G.E. NSFnets (Navier–Stokes flow nets): Physics-informed neural networks for the incompressible Navier–Stokes equations. *J. Comput. Phys.* **2021**, *426*, 109951. [[CrossRef](#)]
39. Khopkar, A.R.; Kasat, G.R.; Pandit, A.B.; Ranade, V.V. CFD simulation of mixing in tall gas–liquid stirred vessels: Role of local flow patterns. *Chem. Eng. Sci.* **2006**, *61*, 2921–2929. [[CrossRef](#)]
40. Nienow, A.W. On impeller circulation and mixing effectiveness in the turbulent flow regime. *Chem. Eng. Sci.* **1997**, *52*, 2557–2565. [[CrossRef](#)]
41. Deglon, D.A.; Meyer, C.J. CFD modelling of stirred tanks: Numerical considerations. *Miner. Eng.* **2006**, *19*, 1059–1068. [[CrossRef](#)]

42. Reid, A.; Rossi, R.; Cottini, C.; Benassi, A. CFD simulation of a Rushton turbine stirred tank using open-source software with critical evaluation of MRF-based rotation modeling. *Meccanica* **2025**, *60*, 1613–1637. [[CrossRef](#)]
43. Patil, H.; Patel, A.K.; Pant, H.J.; Venu Vinod, A. CFD simulation model for mixing tank using multiple reference frame (MRF) impeller rotation. *ISH J. Hydraul. Eng.* **2021**, *27*, 200–209. [[CrossRef](#)]
44. Barros, P.L.; Ein-Mozaffari, F.; Lohi, A. Analysis of scale-up and characterization of the volumetric mass transfer coefficient for gas dispersion in yield-pseudoplastic fluids using a coaxial mixer. *Ind. Eng. Chem. Res.* **2024**, *63*, 11157–11170. [[CrossRef](#)]
45. Lane, G.L.; Schwarz, M.P.; Evans, G.M. Numerical modelling of gas–liquid flow in stirred tanks. *Chem. Eng. Sci.* **2005**, *60*, 2203–2214. [[CrossRef](#)]
46. Laakkonen, M.; Moilanen, P.; Alopaeus, V.; Aittamaa, J. Modelling local bubble size distributions in agitated vessels. *Chem. Eng. Sci.* **2007**, *62*, 721–740. [[CrossRef](#)]
47. Taghavi, M.; Zadghaffari, R.; Moghaddas, J.; Moghaddas, Y. Experimental and CFD investigation of power consumption in a dual Rushton turbine stirred tank. *Chem. Eng. Res. Des.* **2011**, *89*, 280–290. [[CrossRef](#)]
48. Drew, D.A.; Passman, S.L. *Theory of Multicomponent Fluids*; Springer: New York, NY, USA, 2006.
49. Ishii, M.; Hibiki, T. *Thermo-Fluid Dynamics of Two-Phase Flow*; Springer: New York, NY, USA, 2010.
50. Elman, H.C.; Ramage, A.; Silvester, D.J. IFISS: A computational laboratory for investigating incompressible flow problems. *SIAM Rev.* **2014**, *56*, 261–273. [[CrossRef](#)]
51. Fukami, K.; Fukagata, K.; Taira, K. Machine-learning-based spatio-temporal super-resolution reconstruction of turbulent flows. *J. Fluid Mech.* **2021**, *909*, A9. [[CrossRef](#)]
52. Wu, C.; Zhu, M.; Tan, Q.; Kartha, Y.; Lu, L. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *arXiv* **2022**, arXiv:2207.10289.
53. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
54. Yu, J.; Lu, L.; Meng, X.; Karniadakis, G.E. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *Comput. Methods Appl. Mech. Eng.* **2022**, *393*, 114823. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.