processes

MDPI

Article

Cross-Geometry Flow-Field Prediction in Irregular Wellbores Using Parameterized Physics-Informed Neural Networks

Zhichao Xie [1,2,3], Jian Yang [1], Chi Zhao [2,3,*], Long Chen [1], Jiali Wang [1] and Yunhai Zhao [1]

[1] Oil and Gas Technology Research Institute, PetroChina Southwest Oil & Gasfield Company, Chengdu 610017, China; swpu_xzc@163.com (Z.X.); 18180712593@163.com (J.Y.); chenlong01@petrochina.com.cn (L.C.); wangjiali2022@petrochina.com.cn (J.W.); zyh11230715@163.com (Y.Z.)
[2] National Engineering Research Center for Oil & Gas Drilling and Completion Technology, Yangtze University, Wuhan 430100, China
[3] Hubei Key Laboratory of Oil and Gas Drilling and Production Engineering, Yangtze University, Wuhan 430100, China
* Correspondence: 2023730014@yangtzeu.edu.cn

Abstract

Irregular borehole enlargement alters annular flow structures and can create low-velocity regions that impair hole cleaning. However, conventional computational fluid dynamics (CFD) requires repeated geometric reconstruction and numerical simulations for each wellbore configuration, while existing physics-informed models have limited capability for cross-geometry extrapolation. This study develops transfer-learning-based parameterized physics-informed neural networks (TLP-PINNs) for the rapid three-dimensional prediction of Herschel–Bulkley drilling-fluid flow in locally enlarged irregular wellbores. A Gaussian enlargement ratio is introduced as an explicit geometric input, while mass and momentum conservation equations, boundary conditions, adaptive loss weighting, extrapolation-distance-driven physics-point allocation, and partial-layer transfer learning are integrated into a unified framework. The source model generally maintains good in-range interpolation accuracy for the velocity field within the 10–50% enlargement range, although the pressure error shows greater case-to-case variability for unseen intermediate geometries. For the unseen 80% enlargement case, transfer learning reduces the relative RMSEs of axial velocity, velocity magnitude, and pressure from 14.0%, 13.5%, and 3.5% to 9.42%, 9.50%, and 1.34%, respectively, with corresponding R_p^2 of 0.9668, 0.9663, and 0.9992. The transferred model also accurately captures the low-velocity zone within the enlarged section and the associated redistribution of the local pressure gradient, while reducing neural-network training time by 51.2% relative to full-domain retraining. These results demonstrate a practical cross-geometry framework for the rapid hydraulic assessment of irregular wellbores and provide detailed flow-field information for subsequent analyses of hole cleaning and particle transport.

Keywords: irregular wellbore; physics-informed neural networks; parameterized PINNs; transfer learning; Herschel–Bulkley fluid; flow-field prediction

Check for updates

Academic Editor: Evangelos Tsotsas

Received: 3 September 2026
Revised: 26 October 2026
Accepted: 29 October 2026
Published: 3 November 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) license.

1. Introduction

During drilling operations, wellbore instability, mechanical erosion, formation collapse, and lithological heterogeneity can lead to borehole enlargement, ovalization, and other localized geometric irregularities. These geometric variations alter the annular cross-sectional area and consequently redistribute the velocity field, pressure gradient, and wall shear stress.

For yield-stress and shear-thinning drilling fluids, enlarged sections may further promote the formation of low-velocity or nearly stagnant zones, thereby impairing cuttings suspension and transport. Previous experimental and numerical studies have demonstrated that annular geometry, drill-pipe eccentricity and rotation, and fluid rheology strongly affect cuttings transport and pressure loss [1,2]. Therefore, accurate prediction of the velocity and pressure fields in irregular wellbores is essential for characterizing local flow redistribution and identifying regions that may be unfavorable for effective hole cleaning.

The effects of wellbore irregularity on annular flow have been extensively investigated [3,4]. Skadsem et al. [5] combined experiments with three-dimensional computational fluid dynamics (CFD) simulations to examine displacement flow in highly inclined irregular annuli, whereas Etrati et al. [6] investigated drilling-fluid removal from washout regions under different rheological and operating conditions. Renteria et al. [7] further analyzed the effects of borehole ovalization and axial diameter variations in horizontal wells, while related CFD studies have examined flow behavior in irregular wellbores [8]. Collectively, these studies demonstrate that borehole enlargement can substantially alter local velocity and pressure distributions. However, most existing investigations treat individual geometries independently. For each new borehole configuration, conventional CFD generally requires repeated geometric reconstruction, mesh generation, and numerical solution. Consequently, obtaining flow fields over a continuous range of irregular geometries remains computationally expensive, particularly when rapid cross-geometry evaluation is required.

Scientific machine learning provides an alternative approach for parameterized flow-field prediction. Physics-informed neural networks (PINNs) incorporate governing equations and boundary conditions into neural-network training through automatic differentiation [9]. By embedding physical laws into the optimization process, PINNs have been successfully applied to fluid mechanics, heat transfer, multiphase flow, and related parameterized problems [10–14]. Recent advances have further extended physics-informed learning from fixed computational domains to complex geometries and parameterized solution spaces [15–17], providing a potential framework for continuously mapping geometric parameters to corresponding flow fields.

PINNs have also been introduced into wellbore hydraulic and thermal problems, including transient pressure prediction, wellbore thermal analysis, and hydraulic transients [18,19]. Nevertheless, most existing wellbore PINN studies are based on one-dimensional formulations or regular computational domains. For complex-geometry problems, parameterized models are generally more reliable for interpolation within the training range, whereas their predictive accuracy may deteriorate when the target geometry lies outside the source parameter domain. This limitation is particularly relevant to irregular wellbores, in which three-dimensional boundaries, non-Newtonian rheology, and geometric variations are strongly coupled. Transfer learning provides a promising strategy by reusing knowledge acquired under source conditions and adapting selected network representations to related target tasks, and recent studies have demonstrated its effectiveness for physics-informed modeling across varying flow conditions and scales [19,20]. Meanwhile, adaptive loss weighting and physics-point sampling have been shown to improve PINN optimization and physical constraint enforcement [21–25]. However, cross-geometry prediction of three-dimensional non-Newtonian wellbore flow that integrates geometric parameterization, extrapolation assessment, target-domain physical constraints, and transfer learning remains insufficiently explored [26].

To address these limitations, this study proposes a transfer-learning-based parameterized physics-informed neural-network (TLP-PINN) framework for rapid flow-field prediction in irregularly enlarged wellbores. A continuous Gaussian function is first used to describe local borehole enlargement, and the spatial coordinates together with the enlarge-

ment ratio, α, are incorporated into the network input, while the three velocity components and pressure constitute the network outputs. The continuity and momentum equations governing Herschel–Bulkley non-Newtonian flow, together with the inlet and outlet conditions and no-slip wall boundary conditions, are imposed as physical constraints. Unlike conventional approaches that train an independent network for each enlargement ratio, the source model first learns shared spatial and geometric features from cases with relatively small enlargement ratios. Independent cases are then used to assess parameterized interpolation within the training range. Cases with larger enlargement ratios are subsequently treated as target conditions outside the source parameter domain to quantify direct extrapolation errors. Based on the distance between the target condition and the source parameter domain, the distribution of physical collocation points is adaptively adjusted. Cross-geometry transfer is then achieved through source-model initialization, partial-layer freezing, and fine-tuning with a reduced learning rate.

The remainder of this paper is organized as follows. Section 2 introduces the irregular-wellbore geometry, governing equations, boundary conditions, CFD reference solutions, and case design. Section 3 presents the proposed TLP-PINN framework, including the physics-informed loss function, adaptive weighting strategy, physics-point allocation, network architecture, and transfer-learning procedure. Section 4 evaluates parameterized interpolation, direct extrapolation, target-domain transfer, flow-field reconstruction, and computational performance. Finally, Section 5 summarizes the main conclusions and discusses the limitations of the present study and potential directions for future work.

2. Physical Model and Problem Definition

2.1. Geometric Model of the Irregular Wellbore

As shown in Figure 1, a typical wellbore system in modern drilling operations primarily consists of cased sections, the drillstring, the annulus, and a horizontal open-hole section. Drilling fluid is pumped downward through the drillstring to the bit and then returns to the surface through the annulus. Owing to wellbore erosion, formation collapse, and geological heterogeneity, the open-hole section is susceptible to geometric irregularities such as localized borehole enlargement, which alters the flow cross-sectional area and redistributes the velocity, pressure, and wall shear stress fields. As the drilling fluid flows through these irregular sections, continuous variations in cross-sectional area and boundary geometry can further modify the local flow structure and may promote flow separation and localized recirculation [27].

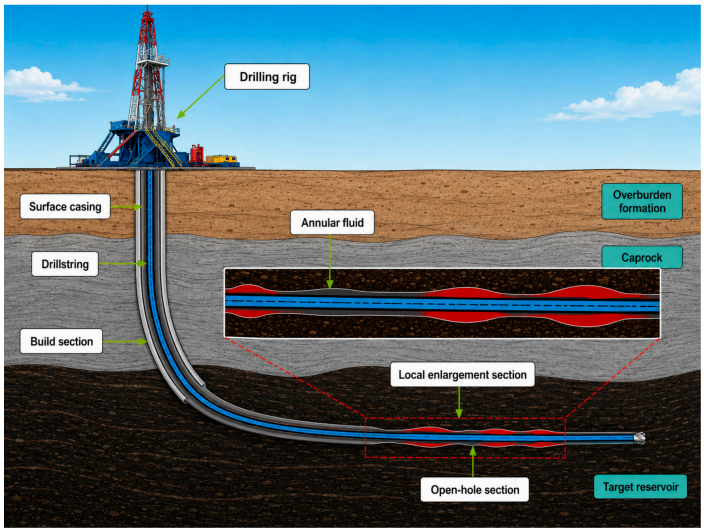


Figure 1. Irregular wellbore schematic diagram.

In this study, the locally irregular wellbore section is simplified as a concentric, axisymmetric annulus with a Gaussian-shaped enlargement. As shown in Figure 2, the drillstring axis is aligned with the (x)-axis, while the (y)- and (z)-axes lie in the cross-sectional plane. Let (r) denote the radial distance from an arbitrary point in the cross section to the drillstring axis. The three-dimensional computational domain is then characterized by an axially varying wellbore radius, which is defined using the following functions:

$$\begin{cases} r = \sqrt{y^2 + z^2}, \\ R_w(x;\alpha) = R_0\left[1 + \alpha \exp\left(-\dfrac{(x - x_c)}{2\sigma_g^2}\right)\right], \\ \Omega(\alpha) = \{(x,y,z) | 0 \le x \le L, R_d \le r \le R_w(x;\alpha)\}, \\ \alpha = \dfrac{R_{\max} - R_0}{R_0} = \dfrac{D_{\max} - D_0}{D_0}, \end{cases} \tag{1}$$

where R_0 and D_0 are the reference wellbore radius and diameter, respectively, m; R_d and D_d are the drillstring radius and outer diameter, respectively, m; α is the dimensionless borehole enlargement ratio, defined as the difference between the maximum borehole diameter and the reference borehole diameter normalized by the reference borehole diameter; x_c denotes the center of the enlarged section, m; σ_g is the local widening width scale along the axial direction of the irregular wellbore, m; x_s and x_e denote the starting and ending positions of the enlarged section, respectively, m; and L is the length of the local computational domain, m. The wellbore radius varies continuously according to the Gaussian function without abrupt truncation, consistent with the geometry used in the CFD simulations. This Gaussian profile provides an idealized representation of smooth, localized borehole enlargement while capturing the associated axial variations in annular geometry. However, more complex geometric features, such as asymmetric washouts, borehole ovalization, and irregular wall roughness, are not explicitly considered. Therefore, the present findings are primarily applicable to the concentric Gaussian-enlargement configurations investigated in this study. Future work will extend the geometric parameterization to abrupt enlargements and more complex irregular wellbore geometries.

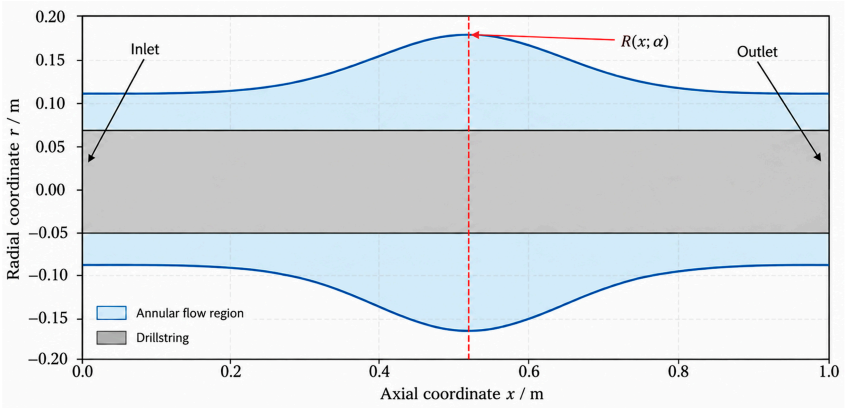


Figure 2. Axial simplified model diagram of the local expansion section.

2.2. Governing Equations and Boundary Conditions

The rheological properties of drilling fluids directly affect annular pressure loss, local velocity distributions, and hole-cleaning performance. In practical drilling operations, drilling fluids generally exhibit pronounced non-Newtonian behavior [28]. Among commonly used constitutive models, the Herschel–Bulkley model can simultaneously represent yield-stress and shear-thinning characteristics and is therefore well suited for modeling annular drilling-fluid flow. In this study, the drilling fluid is treated as a steady, incom-

pressible, isothermal, single-phase Herschel–Bulkley fluid. To isolate the influence of local borehole enlargement on flow-field redistribution, drillstring rotation and gravitational body forces are neglected, and the fluid density is assumed to remain constant [29].

In a Cartesian coordinate system, the velocity vector is denoted by $\mathbf{u} = (u, v, w)$, and the pressure is denoted by p. For a steady, incompressible generalized Newtonian fluid, the mass conservation equation, momentum conservation equation, and rate-of-deformation tensor can be expressed as follows:

$$\begin{cases} \nabla \cdot \boldsymbol{u} = 0, \\ \rho(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} = -\nabla p + \nabla \cdot \left[2\mu_{eff}(\dot{\gamma})\boldsymbol{D}\right], \\ D = \frac{1}{2}\left[\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^T\right], \end{cases} \tag{2}$$

where ρ is the fluid density, kg·m^{-3}; $\boldsymbol{D}$ is the rate-of-deformation tensor; and μ_{eff} is the effective apparent viscosity, Pa·s. The deviatoric stress contribution is directly represented by $\nabla \cdot (2\mu_eff\mathbf{D})$ in the momentum equation. Because the viscosity of the drilling fluid varies with the local shear state, the viscous term in the momentum equation is determined using the Herschel–Bulkley constitutive relation. For the classical Herschel–Bulkley model, the post-yield shear stress can be expressed as [28]:

$$\tau = \tau_0 + K\dot{\gamma}^n \tag{3}$$

and the corresponding apparent viscosity is:

$$\mu_{HB} = \frac{\tau_0}{\dot{\gamma}} + K\dot{\gamma}^{n-1} \tag{4}$$

As $\dot{\gamma}$ approaches zero, the yield-stress contribution causes the apparent viscosity to increase sharply. To avoid numerical singularities in low-shear regions, a regularized shear rate together with an upper bound on the apparent viscosity is adopted in this study, as follows:

$$\dot{\gamma} = \sqrt{2\boldsymbol{D} : \boldsymbol{D} + \dot{\gamma}_r^2} \tag{5}$$

$$\mu_{eff} = \min(\mu_{\max}, \mu_{HB}) \tag{6}$$

where $\dot{\gamma}$ is the regularized shear rate, s^{-1}; $\dot{\gamma}_r$ is the shear-rate regularization parameter, s^{-1}; τ_0 is the yield stress, Pa; K is the consistency index, Pa·s^n; n is the flow behavior index; μ_{HB} is the Herschel–Bulkley apparent viscosity, Pa·s; and $\mu_{\max}$ is the maximum apparent viscosity, Pa·s. This regularization prevents the apparent viscosity from becoming unbounded in low-shear regions while preserving the rheological characteristics of the Herschel–Bulkley model over the normal shear-rate range.

For the locally irregular annular computational domain shown in Figure 2, the inlet, solid walls, and outlet are denoted by Γ_{in}, Γ_{wall}, and Γ_{out}, respectively. A uniform velocity boundary condition is prescribed at the inlet, while no-slip conditions are imposed on both the drillstring surface and the wellbore wall. At the outlet, a reference pressure is specified together with a zero-normal-gradient condition for the velocity. The corresponding boundary conditions are written as follows [30]:

$$\begin{cases} \boldsymbol{u} = (U_{in}, 0, 0), \quad x \in \Gamma_{in}, \\ \boldsymbol{u} = 0, \quad x \in \Gamma_{wall}, \\ p = p_{out}, \partial u/\partial n = 0, \quad x \in \Gamma_{out}, \end{cases} \tag{7}$$

where U_{in} is the mean axial velocity at the inlet, set to 0.6 m·s^{-1} in this study; p_{out} is the reference pressure at the outlet, set to 0; and **n** denotes the outward unit normal vector to the outlet boundary. No-slip boundary conditions are imposed on both the drillstring surface and the wellbore wall; therefore, all three velocity components vanish at the solid boundaries.

2.3. CFD Reference Solution and Mesh-Convergence Analysis

In this study, the reference flow fields were computed using OpenFOAM v13 with the incompressibleFluid solver, which solves the steady, incompressible continuity and momentum equations. The fluid was treated as a laminar generalized Newtonian fluid governed by the regularized Herschel–Bulkley model described in Section 2.2. Pressure–velocity coupling was handled using the SIMPLE algorithm with two non-orthogonal pressure correctors. The convective, gradient and Laplacian terms were discretized using the bounded Gauss upwind, Gauss linear and Gauss linear corrected schemes, respectively. The pressure equation was solved using a preconditioned conjugate-gradient solver with a diagonal incomplete-Cholesky preconditioner, whereas the velocity equations were solved using smoothSolver with a symmetric Gauss–Seidel smoother. Identical discretization schemes and solver controls were applied to all enlargement-ratio cases to ensure that the resulting CFD reference fields were directly comparable.

The apparent viscosity was evaluated using the regularized Herschel–Bulkley relations given in Equations (3)–(7), with the rheological parameters listed in Table 1. The nominal SIMPLE residual targets were 10^{-5} for velocity and 10^{-4} for pressure. Because no single residual measure is sufficient to establish steady-state convergence, each case was advanced for 3200 SIMPLE iterations and the retained solution was assessed using three complementary criteria: equation residuals, global mass conservation and iterative stability of the pressure drop. At the final iteration, the final residuals reported by the linear solvers were below the specified targets. In addition, the relative inlet–outlet mass-flow imbalance was below 0.001%, and the pressure difference between the inlet- and outlet-adjacent cross-sections changed by less than 1.5% between the solutions saved at iterations 2400 and 3200. Collectively, these results indicated that the computed flow fields were sufficiently steady for use as CFD reference data.

It is noted that the convergence was not determined solely from algebraic residuals. In addition, the inlet and outlet mass-flow rates and the total pressure drop across the computational domain were monitored during the iterations. A solution was regarded as steady when all prescribed residual tolerances were satisfied and both the relative mass-flow imbalance and the variation of pressure drop remained below the specified thresholds.

Furthermore, a mesh-independence analysis was performed for the 80% enlargement case, which is the principal target condition used for the subsequent transfer-learning and detailed flow-field analyses. Three successively refined meshes containing 107,520, 184,320, and 360,000 cells were considered. To provide a sensitive assessment of spatial resolution in the enlargement region, the axial distribution of the mean axial-velocity gradient magnitude was compared, as shown in Figure 3. All three meshes reproduce the same overall distribution and the two major gradient peaks associated with the geometric transition regions. The coarse mesh shows noticeable local deviations near the peak-gradient regions, whereas the medium and very fine meshes are nearly coincident over most of the domain, with only minor differences around the local extrema. These results indicate that further mesh refinement has a limited influence on the resolved flow-gradient characteristics. Therefore, the medium mesh with 184,320 cells was adopted for the subsequent CFD reference calculations considering both numerical accuracy and computational cost.

Table 1. Main geometric and flow parameters.

Item	Value	Unit
Wellbore diameter	215.9	mm
Drillstring outer diameter	127.0	mm
Hydraulic diameter	88.9	mm
Computational domain length	1.0	m
Drilling fluid density	1028	$kg \cdot m^{-3}$
Start position of borehole enlargement	0.4	m
End position of borehole enlargement	0.6	m
Gaussian width parameter	0.05	m
Drilling fluid inlet velocity	0.6	$m \cdot s^{-1}$
Rheological model	Herschel–Bulkley	-
Consistency index	1.02	$Pa \cdot s^{n}$
Flow behavior index	0.5	Non-dimensional
Yield stress	2.55	Pa
Regularized shear rate	0.01	s^{-1}
Maximum effective viscosity	300	Pa·s
Equivalent Reynolds number	295	Non-dimensional

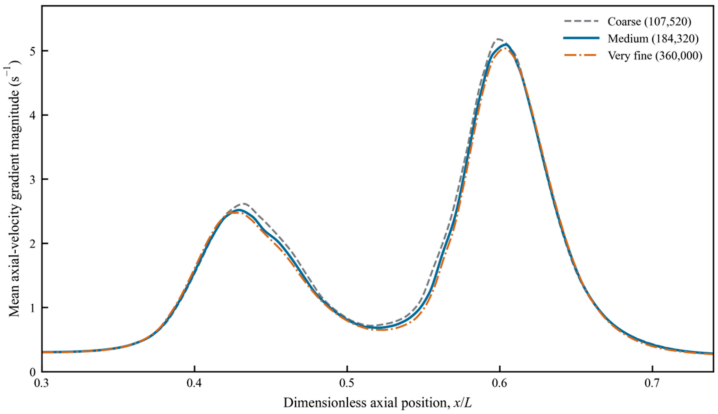


Figure 3. Comparison of the mean axial-velocity gradient magnitude along the axial direction for three mesh resolutions in the 80% enlargement case.

2.4. Case Design and Dataset Partitioning

In this study, the reference flow fields under different borehole enlargement conditions were obtained using OpenFOAM and subsequently used for independent evaluation of the TLP-PINNs model. The computational domain was defined as the locally enlarged three-dimensional annulus shown in Figure 2. Because the present study focuses on flow-field redistribution induced by local borehole enlargement, only a 1 m long local wellbore section was considered in the numerical simulations. The drilling fluid was described as a steady, incompressible Herschel–Bulkley non-Newtonian fluid using the governing equations and regularized constitutive relation introduced in Section 2.2. Accordingly, the reference CFD solutions and the physics residuals embedded in TLP-PINNs were governed by the same continuity and momentum equations and the same effective-viscosity formulation. Except for the borehole enlargement ratio, all other geometric parameters, fluid properties, and boundary conditions were kept identical among the different cases, allowing the influence of borehole enlargement on the local flow field to be examined independently. In this paper, the equivalent Reynolds number based on the apparent viscosity was calculated using the characteristic circumferential shear rate formula ($\dot{\gamma}_c = 8U_{in}/D_h$) proposed by Gulraizet et al. [2], and it was approximately $Re_{eff} \approx 295$. The principal geometric dimensions and flow parameters are summarized in Table 1.

To evaluate the predictive capability of the parameterized model across different regions of the parameter space, all cases were categorized into source-domain training, in-range interpolation, direct extrapolation, and target-domain transfer groups. The source-domain cases were used to train the shared parameterized model, whereas unseen cases within the source training range were used to assess interpolation performance in the continuous parameter space. Cases with larger enlargement ratios were first evaluated by direct extrapolation, without updating the model parameters, to assess the out-of-domain generalization capability of the source model. These cases were subsequently treated as target domains for transfer learning. This design enables a direct comparison of predictive performance before and after transfer under identical target conditions.

Source-domain training incorporated the physics collocation points, with the latter used to enforce the continuity and momentum equations. Independent CFD evaluation datasets were constructed for all interpolation and extrapolation cases and were excluded from source-model training. During target-domain transfer, only a small subset of CFD samples from each target case was used for fine-tuning, while the remaining samples were retained as an independent test set. This separation prevents information leakage between model training and performance evaluation. The case partitioning and corresponding data settings are summarized in Table 2.

Table 2. Case setting and application.

Item	Parameter	Value
Source domain	Enlargement ratio	10%, 30%, 50%
Interpolation	Enlargement ratio	20%, 40%
Direct extrapolation	Enlargement ratio	80%, 100%
Target-domain transfer	Enlargement ratio	80%

For the three source-domain cases with enlargement ratios of 10%, 30%, and 50%, as listed in Table 2, computational meshes were generated in OpenFOAM before model training. Each mesh contained 197,472 points, 566,016 faces, and 184,320 structured hexahedral cells. Mesh quality was assessed using the checkMesh-allTopology-allGeometry utility, and all three meshes successfully passed the quality checks with the message Mesh OK. The corresponding maximum non-orthogonality values were 11.24°, 30.80°, and 44.82°, while the maximum skewness values were 0.109, 0.235, and 0.307, respectively. The maximum aspect ratio was 3.585 for all three meshes. In addition, each mesh consisted of a single connected region and contained no negative-volume cells or non-hexahedral elements.

3. TLP-PINNs

Transfer learning enables a model to adapt to a related target task with limited target-domain data and reduced optimization cost by reusing network parameters and feature representations learned from a source task [27]. In this study, parameterized physics-informed neural networks are integrated with transfer learning to develop a transfer-learning-based parameterized physics-informed neural-network (TLP-PINN) framework. The proposed model enables rapid transfer from a source domain with a low borehole enlargement ratio to a target domain with a higher enlargement ratio, while providing a unified prediction of the velocity and pressure fields under different irregular wellbore geometries. The overall framework of the model is illustrated in Figure 4.

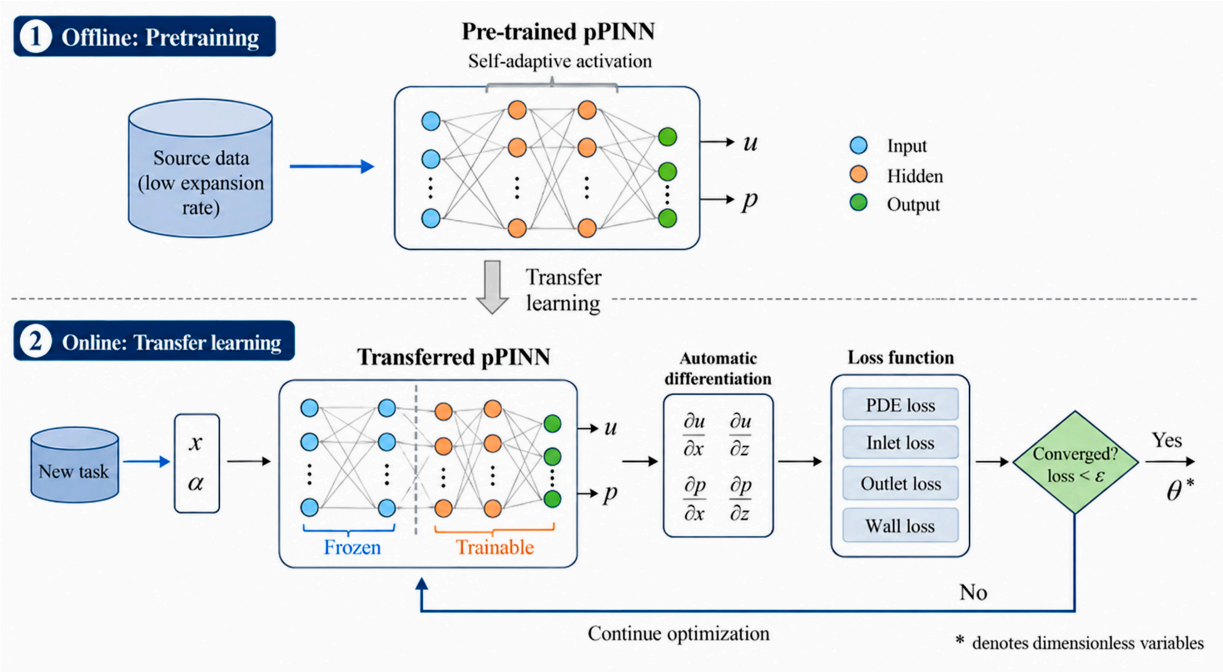


Figure 4. TLP-PINNs model reconstructs the flow-field structure diagram of the expanded section of irregular wellbore.

The network inputs consist of the normalized three-dimensional spatial coordinates and borehole enlargement ratio ($\xi = (\hat{x}, \hat{y}, \hat{z}, \hat{\alpha})$), while the outputs are the three velocity components and pressure ($q = (u, v, w, p)$). By introducing the enlargement ratio α as an explicit conditioning variable, flow fields corresponding to different enlargement conditions can be represented using a shared set of network parameters. Consequently, the learning objective is transformed from a collection of mutually independent flow fields into a joint solution defined over the coupled spatial–geometric parameter domain. The network mapping can therefore be written as:

$$f_\theta : (\hat{x}, \hat{y}, \hat{z}, \hat{\alpha}) \rightarrow (u, v, w, p) \tag{8}$$

To reduce the influence of scale disparities among different physical variables during training, both the spatial coordinates and the enlargement ratio are scaled to a common normalized interval. For numerical scaling, velocity and pressure are normalized using the reference velocity U_{ref} and reference pressure P_{ref}, respectively. During the evaluation of the physics-based residuals, the normalized network outputs are converted to the dimensionless physical variables defined below, and the required spatial derivatives are evaluated using automatic differentiation.

3.1. Physics Constrain Loss

The TLP-PINN is trained under the simultaneous constraints of the governing equations and boundary conditions. For any given borehole enlargement ratio α, the network first predicts the velocity and pressure fields. Automatic differentiation is then employed at the interior collocation points of the computational domain to evaluate the continuity residual and the momentum residuals in the three spatial directions. Consistent with the governing equations introduced in Section 2, these residuals are defined as follows:

$$\begin{cases} F_c = \nabla^* \cdot \boldsymbol{u}^*, \\ \boldsymbol{F}_m = (\boldsymbol{u}^* \cdot \nabla^*)\boldsymbol{u}^* + \nabla^* p^* - \dfrac{1}{\text{Re}_{ref}} \nabla^* \cdot \left[2\mu_{eff}\boldsymbol{D}\right]^* = [F_x, F_y, F_z]^T. \end{cases} \tag{9}$$

where the superscript (*) denotes dimensionless variables. The hydraulic diameter D_h and inlet velocity U_{ref} are adopted as the reference length and velocity, respectively. Accordingly, $\boldsymbol{x}^* = \boldsymbol{x}/D_h$, $\boldsymbol{u}^* = \boldsymbol{u}/U_{ref}$, and $\boldsymbol{p}^* = \boldsymbol{p}/\left(\rho U_{ref}^2\right)$. The dimensionless effective viscosity is defined as $\mu_{eff}^* = \mu_{eff}/\mu_{ref}$, where μ_{ref} is the reference apparent viscosity evaluated from the Herschel–Bulkley model. Accordingly, the reference pressure is $P_{\mathrm{ref}} = \rho U_{\mathrm{ref}}^2$. The dimensional shear rate is obtained from the dimensionless velocity gradients as $\dot{\gamma} = (U_{\mathrm{ref}}/D_h)\sqrt{2\mathbf{D}^* : \mathbf{D}^*}$, where $\mathbf{D}^*$ is the dimensionless rate-of-deformation tensor. The corresponding reference Reynolds number is $Re_{ref} = \rho U_{ref} D_h/\mu_{ref}$. F_c denotes the continuity residual, and $\mathbf{F}_m$ denotes the momentum residual vector, with its three components represented by F_x, F_y, and F_z, respectively. In this study, N_f denotes the number of interior physics collocation points, with 24,000 points used for source-domain pretraining and 6000 points adopted as the baseline for target-domain transfer. The target-domain sampling size is further adjusted according to the extrapolation-distance-driven strategy described in Section 3.3. Accordingly, the total residual loss associated with the governing equations is written as:

$$L_{PDE} = \frac{1}{N_f}\sum_{i=1}^{N_f}\left(F_{c,i}^2 + F_{x,i}^2 + F_{y,i}^2 + F_{z,i}^2\right). \tag{10}$$

For the inlet boundary Γ_{in}, a mean axial velocity of $U_{in} = 0.6\ \mathrm{m{\cdot}s^{-1}}$ is prescribed, while the transverse velocity components are set to zero. The inlet boundary loss is evaluated using $N_{\mathrm{in}} = 900$ collocation points and is defined as:

$$L_{in} = \frac{1}{N_{in}}\sum_{i=1}^{N_{in}}\left[\left(\hat{u}_i - U_{in}\right)^2 + \hat{v}_i^2 + \hat{w}_i^2\right]. \tag{11}$$

No-slip conditions are imposed on both the outer surface of the drillstring and the wellbore wall. A total of $N_{\mathrm{wo}} = 1000$ and $N_{\mathrm{wi}} = 600$ collocation points are sampled on these two wall boundaries, respectively. The corresponding wall-boundary loss is defined as:

$$L_{wall} = \frac{1}{N_{wo}}\sum_{i=1}^{N_{wo}}\left\|\hat{\boldsymbol{u}}_i\right\|^2 + \frac{1}{N_{wi}}\sum_{i=1}^{N_{wi}}\left\|\hat{\boldsymbol{u}}_i\right\|^2. \tag{12}$$

At the outlet boundary, a zero reference pressure and zero normal gradients of the velocity components are imposed. A total of $N_{\mathrm{out}} = 900$ outlet collocation points are used. The outlet-boundary loss is written as:

$$L_{out} = \frac{1}{N_{out}}\sum_{i=1}^{N_{out}}\left\|p_i\right\|^2 + \frac{1}{N_{out}}\sum_{i=1}^{N_{out}}\left\|\frac{\partial \boldsymbol{u}_i}{\partial n}\right\|_2^2, \tag{13}$$

where L_{out} represents the combined outlet-boundary loss, including the outlet pressure and normal velocity gradient.

By combining the governing-equation and boundary-condition constraints described above, the overall loss function employed in this study can be expressed as:

$$L_{total} = \lambda_{\mathrm{PDE}} L_{\mathrm{PDE}} + \lambda_{\mathrm{in}} L_{\mathrm{in}} + \lambda_{\mathrm{wall}} L_{\mathrm{wall}} + \lambda_{\mathrm{out}} L_{\mathrm{out}} \tag{14}$$

where λ denotes the weighting coefficient assigned to each loss component. Its specific definition and calculation procedure are described in Section 3.2.

3.2. Self-Adaptive Loss Weight

The numerical scales and convergence rates of different loss components are not necessarily consistent. To alleviate this imbalance in the multi-objective optimization, trainable adaptive loss weights are introduced, enabling the network to dynamically adjust the relative emphasis placed on the PDE residual and different boundary-condition constraints during training. Based on the adaptive weighting strategy, the loss function is reformulated as a weighted objective:

$$L_{total}(w, \lambda_{pde}, \lambda_{in}, \lambda_{wall}, \lambda_{out}) = L_{pde}(w, \lambda_{pde}) + L_{in}(w, \lambda_{in}) + L_{wall}(w, \lambda_{wall}) + L_{out}(w, \lambda_{out}) \tag{15}$$

where λ denotes the weighting coefficient associated with each loss component. The optimization objective is to minimize the total loss with respect to the network parameters while maximizing it with respect to the adaptive loss weights. This min–max optimization problem can be expressed as:

$$\min_{w} \max_{\lambda_{PDE}, \lambda_{in}, \lambda_{wall}, \lambda_{out}} L_{total}(w, \lambda_{pde}, \lambda_{in}, \lambda_{wall}, \lambda_{out}) \tag{16}$$

Accordingly, the one-step updates of the network parameters and adaptive loss weights can be written in a unified form as [24]:

$$\begin{cases} w^{k+1} = w^k - \eta_w^k \nabla_w L_{total}(w^k, \lambda_{pde}^k, \lambda_{in}^k, \lambda_{wall}^k, \lambda_{out}^k) \\ \lambda_{pde}^{k+1} = \lambda_{pde}^k + \eta_{pde}^k \nabla_{\lambda_{pde}} L_{total}(w^k, \lambda_{pde}^k, \lambda_{in}^k, \lambda_{wall}^k, \lambda_{out}^k) \\ \lambda_{in}^{k+1} = \lambda_{in}^k + \eta_{in}^k \nabla_{\lambda_{in}} L_{total}(w^k, \lambda_{pde}^k, \lambda_{in}^k, \lambda_{wall}^k, \lambda_{out}^k) \\ \lambda_{wall}^{k+1} = \lambda_{wall}^k + \eta_{wall}^k \nabla_{\lambda_{wall}} L_{total}(w^k, \lambda_{pde}^k, \lambda_{in}^k, \lambda_{wall}^k, \lambda_{out}^k) \\ \lambda_{out}^{k+1} = \lambda_{out}^k + \eta_{out}^k \nabla_{\lambda_{out}} L_{total}(w^k, \lambda_{pde}^k, \lambda_{in}^k, \lambda_{wall}^k, \lambda_{out}^k) \end{cases} \tag{17}$$

where η_w and η_j, with $j \in \{pde, \text{in}, \text{wall}, \text{out}\}$, denote the learning rates for the network parameters and the corresponding adaptive loss weights, respectively.

3.3. Extrapolation-Distance Driven Adaptive Physics Point Allocation

Parameterized PINNs generally exhibit good interpolation capability within the source domain. However, when the target borehole enlargement ratio lies far outside the source-domain range, geometric variations may lead to more pronounced redistribution of local velocity gradients, pressure gradients, and near-wall flow structures. Therefore, the same number of physics collocation points is not used for all target conditions. Instead, the sampling density of governing-equation and boundary-condition points is increased according to the distance between the target parameter and the source domain. This strategy is conceptually consistent with residual-based adaptive sampling, but the increment adopted here is directly driven by the extrapolation distance in parameter space, allowing the required sampling size to be determined in advance before transfer learning.

Let the source-domain enlargement-ratio interval be $[\alpha_{\min}, \alpha_{\max}]$. The normalized distance δ_α from the target condition to the source domain and the corresponding risk amplification factor κ_α are defined as [31]:

$$\delta_\alpha = \frac{\max(\alpha_{\min} - \alpha, 0, \alpha - \alpha_{\max})}{\alpha_{\max} - \alpha_{\min}}. \tag{18}$$

$$k_\alpha = \min[k_{\max}, 1 + 3\delta_\alpha] \tag{19}$$

The number of governing-equation collocation points is scaled by κ_α. The maximum amplification factor is set to 2.0 for wall-boundary points and 1.5 for inlet and outlet

points, thereby concentrating additional physics constraints in target regions with greater extrapolation risk:

$$\begin{cases} N_f = round(k_\alpha N_f^0), \\ N_{wall}^t = round[\min(k_\alpha, 2) N_{wall}^0], \\ N_{in/out}^t = round[\min(k_\alpha, 1.5) N_{in/out}^0]. \end{cases} \quad (20)$$

3.4. Network Architecture and Training Strategy

The TLP-PINNs employ a fully connected feedforward neural network. All computations are performed in double precision to reduce round-off errors associated with high-order automatic differentiation and the evaluation of non-Newtonian viscosity. A two-stage optimization strategy combining Adam and L-BFGS is adopted. For source-domain pretraining, the network is first optimized using Adam with an initial learning rate of 6.0×10^{-4}, which is gradually reduced using a cosine-annealing schedule. This is followed by full-batch refinement using L-BFGS. For transfer to the target domain, Adam fine-tuning is performed with a lower learning rate of 1.0×10^{-4}, followed by further optimization using L-BFGS. To evaluate the sensitivity of the model to parameter initialization, independent training runs are conducted using random seeds from 2026 to 2030, and the final prediction is obtained by averaging the outputs of the resulting ensemble of models. The main training and transfer-learning settings are summarized in Table 3.

Table 3. Main training and transfer-learning settings of TLP-PINNs.

Parameter	Value
Network input	$\hat{x}, \hat{y}, \hat{z}, \hat{\alpha}$,
Network output	u, v, w, p
Network architecture	[4, 160, 160, 160, 160, 160, 4]
Activation/initialization	Tanh/Xavier
Numerical precision	Float64
Source-domain physics points	24,000
Target-domain CFD points	1500
Baseline target-domain PDE points	6000
Target-domain inlet/outlet points	900/900
Target-domain wall points	Outer wall:1000/Inner wall:600
Source-domain optimizer	Adam and L-BFGS
Initial source-domain learning rate	6.0×10^{-4}
Target-domain optimizer	Adam:3000 and L-BFGS:1000
Target-domain learning rate	1.0×10^{-4}
Frozen layers	First three linear layers
Random seeds	2026, 2027, 2028, 2029, 2030

The source model is first jointly pretrained on the 10%, 30%, and 50% borehole-enlargement cases to learn transferable flow features shared across locally enlarged wellbore geometries, including velocity distributions, pressure gradients, and near-wall flow characteristics. For the high-enlargement target cases of 80% and 100%, the pretrained source model is first applied directly, without updating its parameters, to generate extrapolative predictions that serve as the pre-transfer baseline. The target model is subsequently initialized with the pretrained source-model parameters and fine-tuned using a limited amount of target-domain CFD data together with newly generated physics collocation points.

Based on the transfer-learning assumption that earlier network layers capture relatively general spatial and geometric representations, whereas layers closer to the output are more responsive to target-specific flow-field redistribution, the first three linear layers are frozen during transfer learning. Only the remaining hidden layers and the output layer are

updated during target-domain fine-tuning. The overall parameterized adaptive network and target-domain transfer procedure are illustrated in Figure 5.

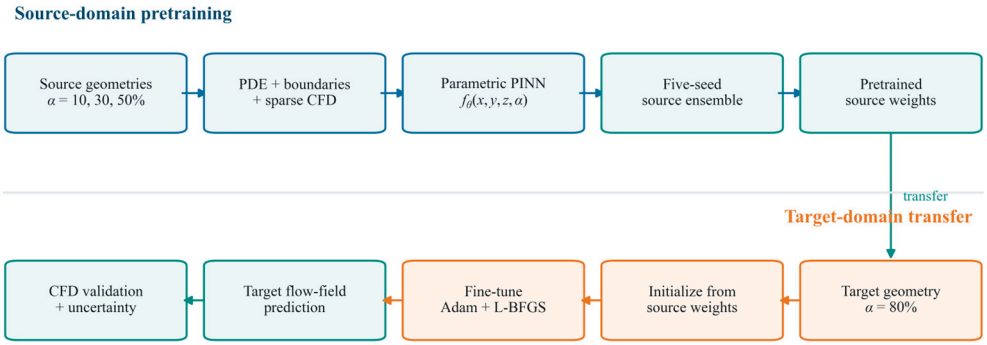


Figure 5. The parametric adaptive network and the target domain transfer process.

In this paper, all computations have been carried out exploiting a workstation equipped with a 12th Inter Core i7-12700KF CPU with 24 GB and an NVIDIA RTX 4060Ti GPU. The source-domain, target-domain transfer, and full retraining cases used the same convergence criteria and stopping strategy to ensure a consistent comparison of computational cost.

4. Results and Discussion

This section evaluates the performance of TLP-PINNs against independent CFD flow fields obtained using OpenFOAM from six perspectives: training stability, in-domain generalization within the source-parameter range, direct extrapolation beyond the training domain, target-domain transfer learning, local flow-field reconstruction, and model comparison [30]. The analysis considers not only global statistical metrics such as RMSE and R^2, but also loss decomposition, scatter distributions, axial and cross-sectional flow fields, and spatial error maps to assess whether the model can accurately recover the local flow structures induced by borehole enlargement.

4.1. Evaluation Metric

To quantitatively assess the differences between the CFD reference solutions and the TLP-PINN predictions, the relative root mean square error (RMSE) and squared Pearson correlation coefficient (R_p^2) are used in this study. The RMSE measures the relative prediction error, whereas R_p^2 evaluates the squared linear correlation between the predicted and reference values. The formulas are given as follows:

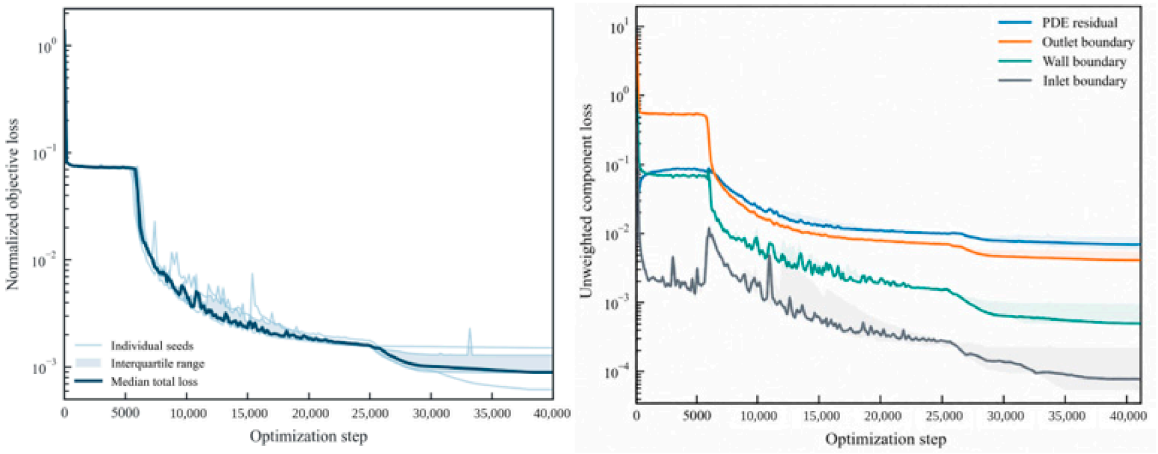
$$Relative\ RMSE = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{q}_i - q_i)^2}}{\sqrt{\frac{1}{N}\sum_{i=1}^{N}q_i{}^2}}, \tag{21}$$

$$R^2 = \frac{\left[\sum_{i=1}^{N}(q_i - \overline{q})(\hat{q}_i - \overline{\hat{q}}_i)\right]^2}{\sum_{i=1}^{N}(q_i - \overline{q})^2\sum_{i=1}^{N}(\hat{q}_i - \overline{\hat{q}}_i)^2}, \tag{22}$$

where q_i and $\hat{q}_i$ denote the CFD reference value and the corresponding model prediction at the i-th evaluation point, respectively; $\overline{q}$ and $\overline{\hat{q}}$ denote the mean values of the CFD reference solution and model prediction, respectively; and N is the total number of evaluation points.

4.2. Training Stability and Generalization of Source-Domain Parameterized Model

The training stability of the source parameterized model under different random initializations is first evaluated to assess whether the observed convergence behavior is robust to initialization. As shown in Figure 6, the total loss decreases consistently across all five random seeds as training progresses. Variations among the different runs are most pronounced during the early stage of training, whereas the loss curves gradually converge toward similar levels and exhibit stable behavior at later stages. The mean final total loss across the five independent runs is 1.103×10^{-4}. Among the individual loss components, both the PDE residual loss decrease consistently, with no evidence of premature stagnation in either component while the total loss continues to decrease. Chen et al. pointed out that instability in PINN training is often associated with imbalanced convergence rates among different residual terms [32]. The multi-seed convergence behavior observed here is consistent with their conclusion that different physical constraints should converge in a coordinated manner. Unlike standard PDE benchmark problems, however, the present study considers three-dimensional non-Newtonian annular flow with a continuous geometric parameter. The convergence of the loss curves toward similar final levels across different random seeds provides direct evidence that the model is relatively insensitive to initialization and further rules out random-seed effects as the dominant source of error in the subsequent interpolation and transfer-learning analyses.



(**a**) Total loss (**b**) Component losses

Figure 6. Training losses of the source parameterized model under multiple random seeds.

Figure 7 further examines whether the different physical constraints remain reasonably balanced during optimization. In terms of their final weighted contributions, the PDE residual accounts for approximately 50.3% of the total objective, the wall-boundary constraint contributes approximately 28.8%, and the inlet and outlet boundary terms together contribute approximately 20.9%. These results indicate that the optimization is primarily governed by the interior conservation equations and the no-slip wall condition rather than by any single boundary term. This behavior is consistent with the findings of Gao et al. and Chen et al. [32,33], who showed that adaptive weighting can alleviate imbalances among multiple loss components. Compared with analyses based solely on the total loss, Figure 7 provides a quantitative decomposition of the weighted contributions of the individual constraints to the final objective function, showing that multiple physical constraints remain active during optimization. The relatively large contribution of the PDE residual is also consistent with its slower decrease during the later stage of training, as shown in Figure 6. Enforcing the three-dimensional conservation equations together with the nonlinear Herschel–Bulkley rheological relation presents a more complex optimization

task than satisfying an individual boundary condition. Consequently, the PDE residual retains a relatively large contribution during the later stages of training.

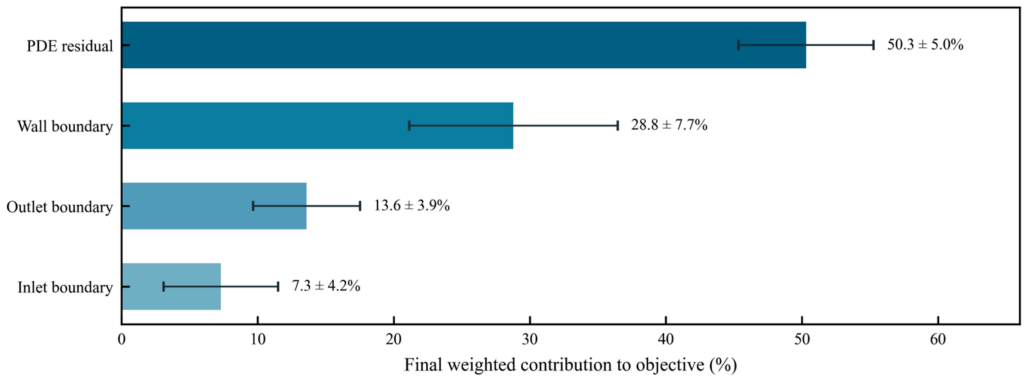
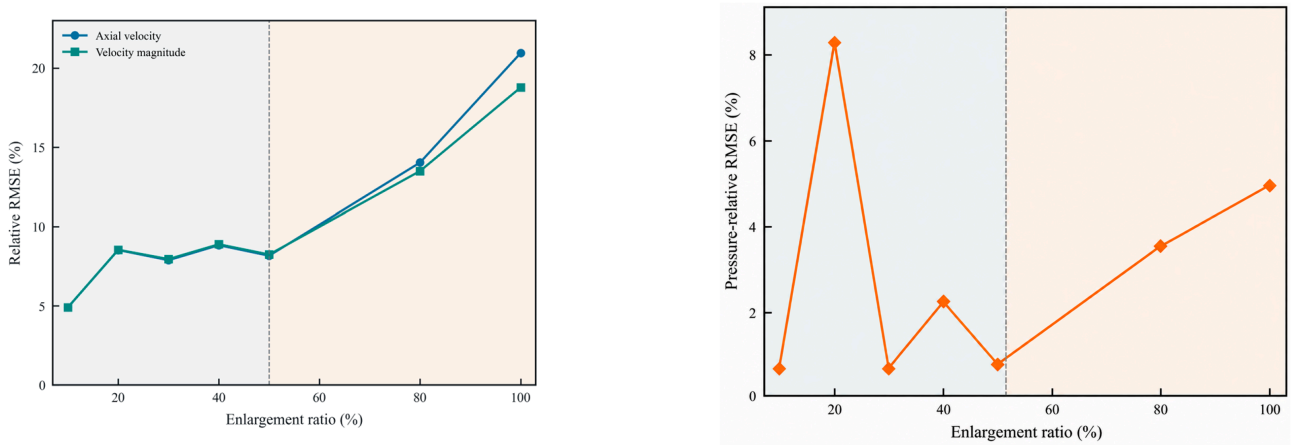


Figure 7. Weighted contributions of individual loss components in the source-domain parameterized model.

4.3. Direct Extrapolation Performance and Transfer Learning Gains

To identify the effective parametric range of the source model, the parameterized network trained on borehole enlargement ratios of 10–50% is first directly applied to larger enlargement ratios without updating the network parameters. As shown in Figure 8, the relative RMSEs of the axial velocity and velocity magnitude remain generally within approximately 5–9% for the source-domain and in-range interpolation cases, but increase markedly once the enlargement ratio exceeds the upper training bound of 50%. At an enlargement ratio of 80%, the relative RMSEs of the axial velocity, velocity magnitude, and pressure are approximately 14.0%, 13.5%, and 3.5%, respectively. For the 100% case, the relative RMSE of the velocity magnitude further increases to 18.78%, while the corresponding R_p^2 decreases to 0.8805. The pressure-relative RMSE varies non-monotonically within the 10–50% range, from approximately 0.6% to 8.3%, and therefore shows no consistent dependence on the enlargement ratio within the source domain. Beyond the 50% training limit, the pressure-relative RMSE increases from approximately 3.5% at 80% to 4.9% at 100%, indicating a deterioration in pressure prediction as the extrapolation distance increases. The relatively large pressure error for the unseen 20% case indicates that in-range interpolation is not uniformly accurate for the pressure field. Compared with velocity, the pressure prediction is influenced by the reconstruction of the global pressure-gradient distribution while being referenced to the prescribed outlet pressure. In addition, because the pressure error is reported in relative form, its magnitude is sensitive to the overall scale of the corresponding CFD pressure field. Consequently, the pressure rRMSE does not necessarily vary monotonically with the enlargement ratio, even within the source-parameter interval. Takao and Ii reported that, in PINN fine-tuning for cavity flows with different geometries, the direct applicability of a pretrained model decreases as the geometric discrepancy between the source and target domains increases [34]. Zhu et al. likewise showed that geometric variations in eccentric configurations can reshape local high-gradient regions and increase the difficulty of transfer learning [35]. In contrast to these studies, the present work establishes a continuous error–enlargement-ratio relationship rather than comparing only a few discrete geometries. The pronounced increase in error beyond the 50% training limit demonstrates that low source-domain losses and a high interpolation R_p^2 cannot substitute for out-of-domain cross-geometry validation. This result also provides the basis for selecting the 80% enlargement case as a representative target for subsequent transfer learning.



(**a**) Relative RMSE of axial velocity and velocity magnitude (**b**) Relative RMSE of pressure

Figure 8. Variation in the relative prediction error of the parameterized model with borehole enlargement ratio. The blue and orange shaded regions denote the training/interpolation range and extrapolation range, respectively.

For the 80% enlargement target domain, the transfer process is further examined to determine whether the data, governing-equation, and boundary-condition constraints can be effectively rebalanced using only a limited amount of target-domain data. During the first 3000 Adam iterations, the total loss and the PDE, inlet, wall, and outlet losses decrease continuously. After switching to L-BFGS refinement, all loss components decrease further in a coordinated manner, indicating that transfer learning not only reduces the data-fitting error using a limited number of CFD samples but also restores physical consistency in the target geometry. This behavior is consistent with the findings of Yan et al. [19], who demonstrated that rapid adaptation to a target task can be achieved by freezing part of a pretrained network while updating the remaining parameters. In the present study, however, the target enlargement ratio increases from the source-domain upper bound of 50% to 80%, while physics collocation points are regenerated for the target geometry and their density is increased according to the extrapolation distance. As shown in Figure 9, both the Adam and L-BFGS stages exhibit stable convergence, and the boundary-condition losses do not increase noticeably after transfer. These results indicate that the pretrained source-model parameters remain transferable and that the physical constraints can be effectively re-established in the target domain.

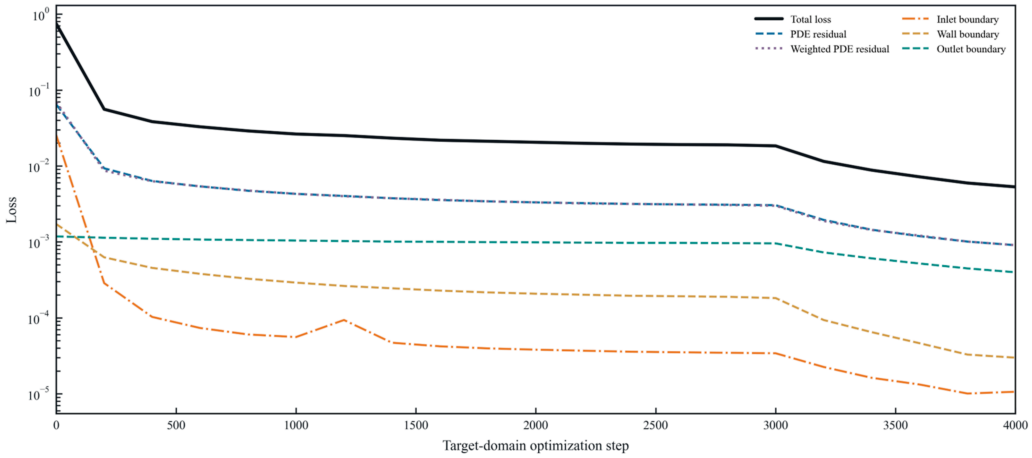
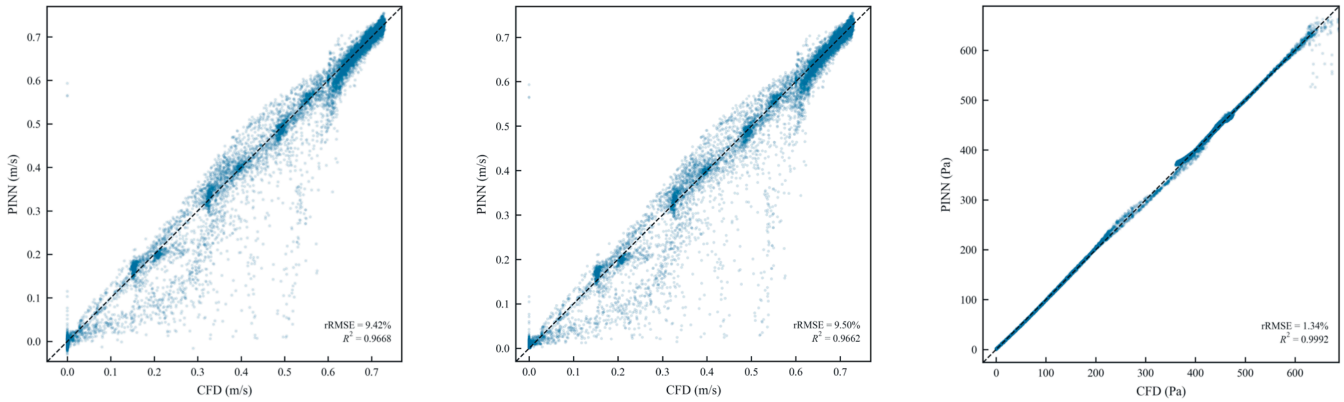


Figure 9. Transfer-training losses of the parameterized model in the 80% enlargement target domain.

Figure 10 further evaluates the full-field consistency of the transferred model based on pointwise prediction distributions. For the 80% enlargement target domain, the relative RMSEs of the axial velocity, velocity magnitude, and static pressure are 9.42%, 9.50%, and 1.34%, respectively, with corresponding R_p^2 values of 0.9668, 0.9663, and 0.9992. The pressure predictions are almost perfectly aligned with the 45° line, whereas the velocity predictions show greater dispersion in the low- and intermediate-velocity ranges. This indicates that the transferred model accurately recovers the overall pressure field, while local velocity variations near the enlargement center remain a major source of error. The remaining discrepancies in the low- and intermediate-velocity regions may be attributed to complex local shear-rate variations and the nonlinear viscosity response of the Herschel–Bulkley fluid. These effects are particularly pronounced near the enlarged section and its transition regions, where limited target-domain samples may not fully resolve the local flow variations. Liu et al. similarly reported that pressure fields are generally easier to predict accurately than local velocity structures in PINN-based transfer learning for complex three-dimensional geometries [36]. This may be attributed to the global pressure constraint imposed by incompressible flow, whereas velocity predictions are more sensitive to local geometric variations and near-wall gradients. Compared with that study, the 80% enlargement case considered here lies entirely outside the geometric range of the source domain and involves non-Newtonian annular flow. Therefore, Figure 10 not only provides overall error metrics but also shows that the dominant prediction errors occur mainly in the low- and intermediate-velocity regions. This observation is further examined using the subsequent spatial error maps.



(**a**) Axial velocity (**b**) Velocity magnitude (**c**) Static pressure

Figure 10. Scatter comparison between CFD results and TLP-PINN predictions for the 80% enlargement case.

Table 4 also presents the evaluation results based on independent CFD sampling points in the 80% enlargement target domain. A lower relative RMSE and an R_p^2 closer to 1 indicate better overall agreement. Since the magnitudes of the transverse velocity components are much smaller than that of the axial velocity, the present discussion of transfer performance focuses primarily on the axial velocity, velocity magnitude, and pressure as the main evaluation variables.

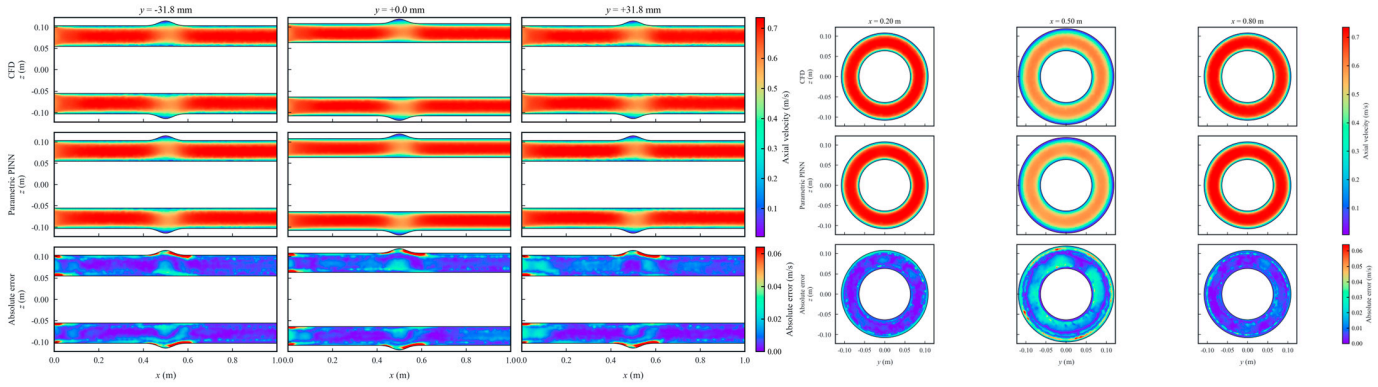
Table 4. Independent evaluation metrics of the transferred model for the 80% enlargement case.

Variable	Relative RMSE	R_p^2
Axial velocity	0.0942	0.9668
Velocity magnitude	0.0950	0.9663
Static pressure	0.0134	0.9992

4.4. Analysis of Flow Field Reconstruction Results

To more clearly demonstrate the flow-field reconstruction capability of the TLP-PINNs model, this section presents the predicted flow fields of the source-domain parameterized PINN model and the TLP-PINNs model for the 10% enlargement case, the 20% enlargement case, and the extrapolated 80% enlargement case. On the other hand, global error metrics alone cannot determine whether the network has truly recovered the local flow structures within the enlarged section. Therefore, the 10% enlargement case, which belongs to the source-domain training conditions, is first selected for a pointwise comparison between CFD and TLP-PINNs on the axial longitudinal section and on cross-sections at different axial locations.

As shown in Figure 11, at the enlargement center ($x \approx 0.50$ m), the increased annular cross-sectional area leads to a pronounced reduction in axial velocity. After the flow leaves the enlarged region, the velocity gradually recovers to the level observed in the regular wellbore section. The model consistently reconstructs the location, radial extent, and axial evolution of the low-velocity region throughout the expansion and contraction process. The absolute error is mainly concentrated in transition regions with large geometric gradients, particularly near the inlet and outlet of the enlarged section. This behavior is consistent with previous studies of irregular wellbores and washout enlargements, which showed that borehole enlargement reduces local flow velocity and may generate low-shear or stagnant regions [37]. Unlike previous studies based primarily on case-specific CFD simulations or experiments, the present study reconstructs the three-dimensional velocity field using a single parameterized network. Figure 11 further provides longitudinal, cross-sectional, and absolute-error distributions, demonstrating that the low-velocity region can be explicitly localized rather than inferred solely from global velocity metrics. This provides a direct flow-field basis for identifying regions potentially unfavorable for cuttings transport in enlarged irregular wellbores.



(**a**) (**b**)

Figure 11. Comparison of axial-velocity fields between CFD and parameterized model predictions for the 10% enlargement case. (**a**) Longitudinal axial-velocity distributions at different transverse positions. (**b**) Cross-sectional axial-velocity distributions at different axial positions.

To evaluate whether this spatial reconstruction capability generalizes to an unseen geometry, Figure 12 presents the interpolation case with a 20% borehole enlargement ratio. Although this case is not included in the supervised training data, the TLP-PINN successfully captures the axial deceleration caused by the increased cross-sectional area near $x \approx 0.50$ m and reproduce the higher velocity distribution of the regular annulus at $x = 0.20$ and 0.80 m. Prediction errors remain concentrated mainly near the geometric transition boundaries and in regions with large radial gradients at the central cross-section. This observation is consistent with Liu et al. [36], who showed that explicitly incorporating

geometric parameters into the network can support continuous prediction across related three-dimensional geometries. The present study further provides sectional evidence for an unseen geometry rather than relying solely on global error metrics.

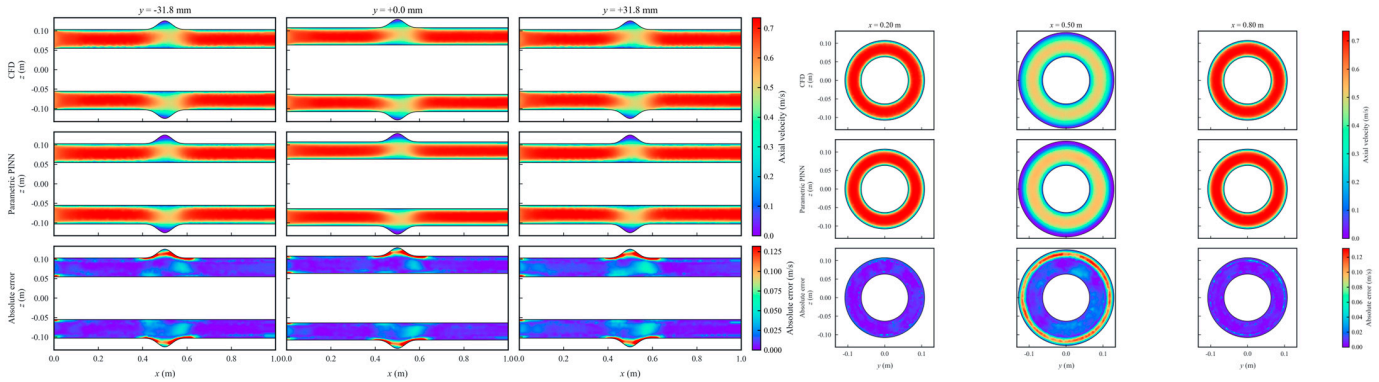


Figure 12. Comparison of axial-velocity fields between CFD and parameterized model predictions for the 20% enlargement case. (**a**) Longitudinal axial-velocity distributions at different transverse positions. (**b**) Cross-sectional axial-velocity distributions at different axial positions.

Compared with the 10% training case, the error band near the center of the enlarged section becomes slightly more pronounced for the 20% case. This may result from differences in the local annular geometry and velocity gradients relative to the discrete training cases. Nevertheless, the error remains localized rather than propagating throughout the computational domain. Figure 12 therefore provides spatial evidence complementing the results in Section 4.1, showing that interpolation within the source-domain parameter range preserves the geometry-dependent flow response rather than merely producing low global errors.

Figure 13 further illustrates why the source model performs poorly under direct extrapolation for the 80% enlargement case and where transfer learning provides improvements. The CFD results show that the large enlargement produces a broader low-velocity region and substantially increases the velocity gradients at the inlet and outlet of the enlarged section. Without transfer learning, the source model captures the approximate location of the low-velocity region but exhibits systematic deviations in its radial extent and gradient magnitude, resulting in pronounced annular or banded high-error regions near the enlargement center and geometric transition boundaries. After transfer learning, these high-error regions shrink markedly, and the radial variation of the cross-sectional velocity distribution agrees more closely with the CFD results. In this study, the model is extrapolated directly from the source-domain upper limit of 50% to an 80% enlargement ratio, while the working fluid exhibits Herschel–Bulkley shear-thinning and yield-stress behavior. The side-by-side comparison of the direct-extrapolation and post-transfer error maps on identical slices shows that the improvement from transfer learning is concentrated primarily in the enlargement transition regions and along the boundaries of the low-velocity zone rather than being uniformly distributed throughout the flow field.

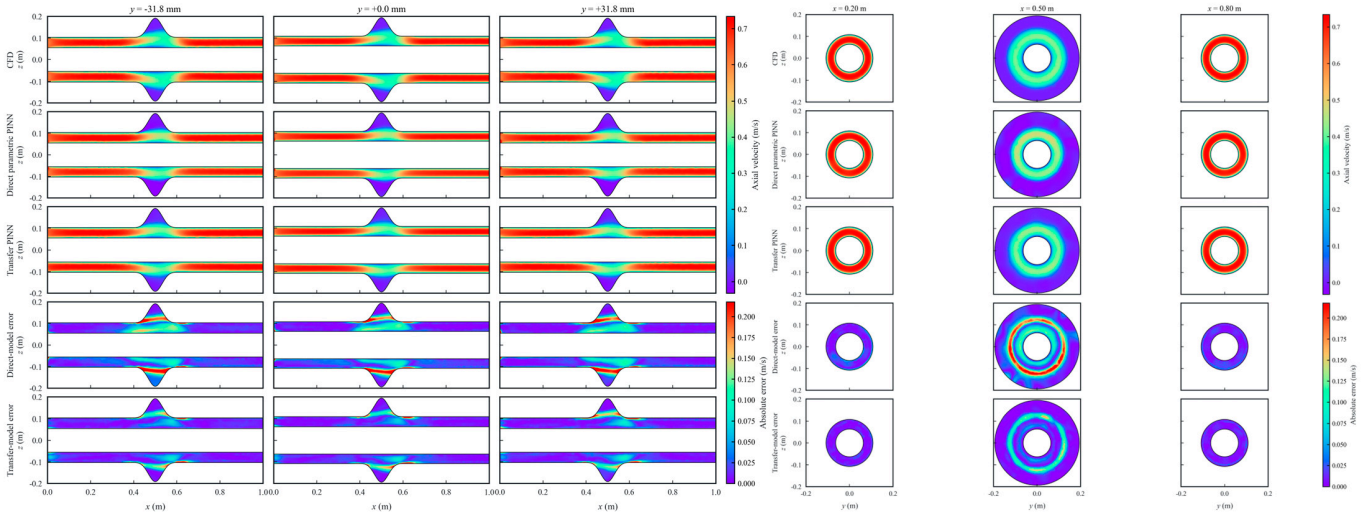
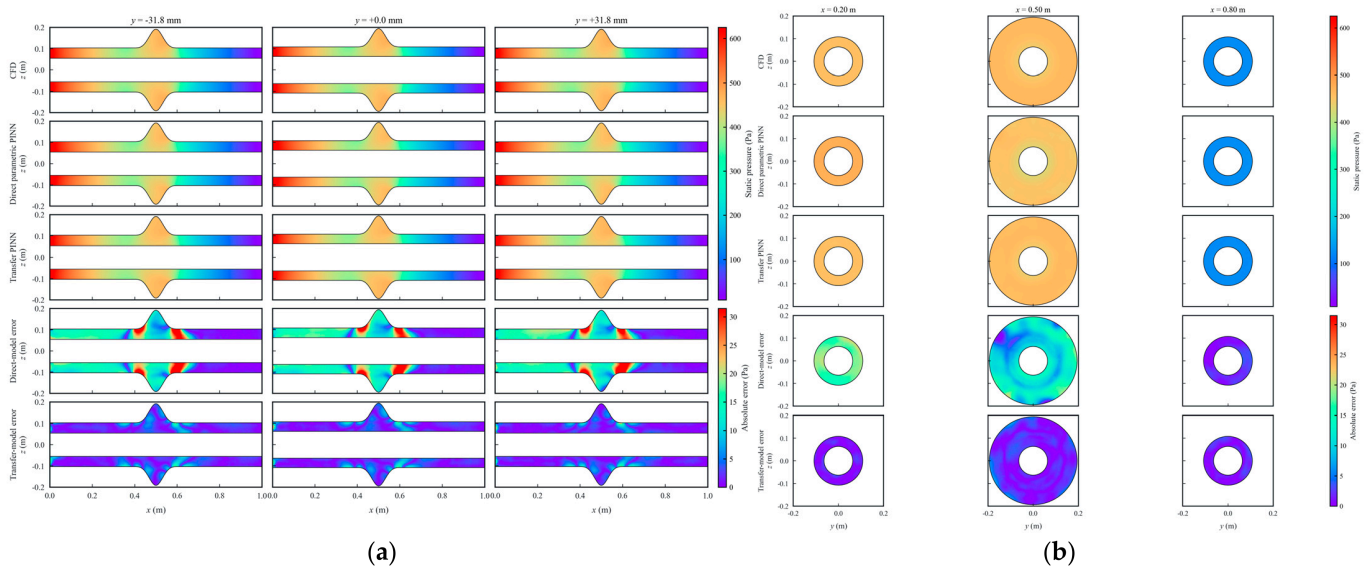


Figure 13. Comparison of axial-velocity slices among CFD, direct extrapolation, and transfer-learning predictions for the unseen 80% enlargement case. (**a**) Longitudinal axial velocity distributions at different trans-verse positions. (**b**) Cross-sectional axial-velocity distributions at different axial positions.

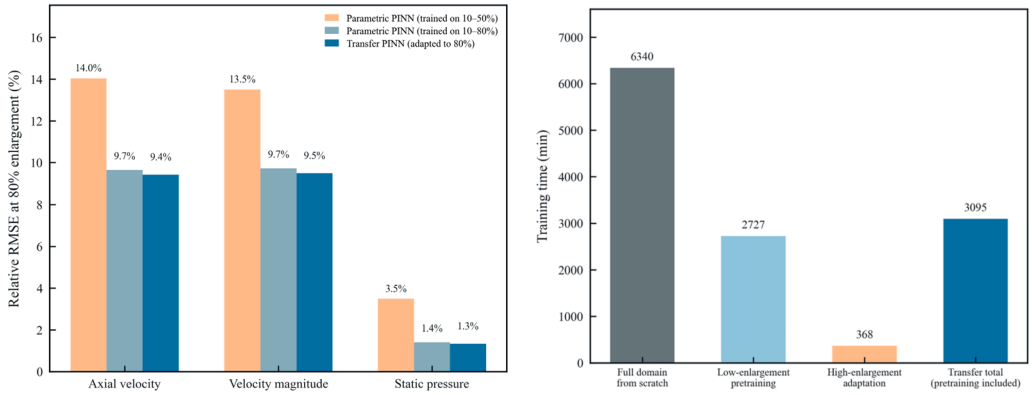
Figure 14 further examines the reconstruction of the pressure field coupled with the velocity field. The CFD pressure decreases smoothly along the axial direction. Although the enlarged section modifies the local pressure gradient, it does not produce spatial variations as sharp as those observed in the velocity field. Direct extrapolation already reproduces the overall pressure level reasonably well, resulting in a substantially lower relative error than that of the velocity field. Nevertheless, localized high-error regions remain near the inlet, center, and outlet of the enlarged section. After transfer learning, both the magnitude and spatial extent of these errors are further reduced, particularly on the central cross-section of the enlargement. Liu et al. also reported high pressure-field reconstruction accuracy using PINN-based transfer learning for complex three-dimensional transport problems [36]. The pressure-relative RMSE of 1.34% and R_p^2 of 0.9992 obtained in the present study are consistent with this trend. In the present case, however, the pressure variation is induced by geometric expansion and contraction in a non-Newtonian annulus, while the target geometry lies outside the training range. Taken together, Figures 13 and 14 demonstrate that the transferred model improves both local velocity-gradient structures and the global pressure distribution rather than optimizing only a single output variable. The reconstructed velocity and pressure fields can further provide inputs for evaluating pressure forces, drag, and lift on particles. However, these results are based on steady single-phase flow and should not be interpreted as direct simulations of cuttings concentration, particle collisions, or cuttings-bed evolution.

Figure 15 evaluates the practical value of transfer learning from the combined perspectives of predictive accuracy and training cost. When the source model trained over the 10–50% enlargement range is directly extrapolated to the 80% case, the relative RMSEs of axial velocity, velocity magnitude, and pressure are 14.0%, 13.5%, and 3.5%, respectively. When the parameterized model is instead expanded and retrained over the full 10–80% range, the corresponding errors decrease to 9.7%, 9.7%, and 1.4%. After transfer learning with TLP-PINNs, the errors are further reduced to 9.4%, 9.5%, and 1.3%, respectively. Thus, without retraining the model over the entire parameter range, the transferred model achieves accuracy comparable to, and slightly better than, that of the fully retrained parameterized model.



(**a**) (**b**)

Figure 14. Comparison of static pressure slices among CFD, direct extrapolation, and transfer-learning predictions for the unseen 80% enlargement case. (**a**) Longitudinal static pressure distributions at different trans-verse positions. (**b**) Cross-sectional static pressure distributions at different axial positions.



(**a**) Accuracy comparison (**b**) Training-time comparison

Figure 15. Transfer-learning gains for the 80% enlargement target domain: (**a**) accuracy comparison; (**b**) training-time comparison.

4.5. Sensitivity to Target-Domain Sample Size

To examine the sensitivity of transfer performance to the amount of target-domain CFD supervision, a sensitivity analysis was conducted for the 80% enlargement case. Notably, the transfer performance does not improve monotonically relative to the direct-extrapolation baseline when only a very small number of target-domain samples is available. With 250 target samples, the velocity RMSE increases to 22.35%, compared with 13.50% for direct extrapolation, indicating a negative-transfer effect under insufficient target-domain information. In this case, the limited target information is insufficient to reliably correct the geometry-dependent flow redistribution, while fine-tuning may perturb some of the transferable representations inherited from the source model. When the number of target samples is increased to 750, the velocity RMSE decreases to 13.42%, which is essentially comparable to the direct-extrapolation result and therefore does not represent a meaningful transfer gain. A clear improvement first appears at 1500 samples, for which the velocity RMSE decreases to 9.50%. Further increasing the number of samples to 3000 and 5000 provides only modest additional improvement. Therefore, 1500 samples represent a reason-

able balance between prediction accuracy and target-domain data requirements. Increasing the sample size from 1500 to 5000 reduces the velocity-magnitude RMSE only from 9.50% to 8.63%, while requiring considerably more CFD data. Accordingly, 1500 samples were selected as a cost-effective setting among the sample sizes investigated. Shown as Table 5.

Table 5. Sensitivity of prediction accuracy to target-domain CFD sample size for the 80% enlargement case.

Target Samples	Velocity rRMSE	Pressure rRMSE
250	0.2235	0.02154
750	0.1342	0.01583
1500	0.0950	0.01341
3000	0.0892	0.01277
5000	0.0863	0.01230

The training-time comparison distinguishes offline source-domain pretraining from target-domain adaptation. Full-domain retraining from random initialization requires 6340 min, whereas source-domain pretraining requires 2727 min and subsequent adaptation to the 80% target domain requires 368 min. Including both stages, the total TLP-PINN training time is 3095 min, representing a 51.2% reduction compared with full-domain retraining. If the source model is already available as a reusable offline model, a single target-domain adaptation requires only approximately 5.8% of the time needed for full-domain training from scratch, corresponding to an approximately 17.2-fold reduction in training time. Compared with these previous studies [35], the efficiency assessment adopted here is more conservative because it reports not only the target-domain fine-tuning time but also the source-domain pretraining cost as part of the "transfer total." Absolute acceleration factors should not be compared directly across different studies because of differences in geometry, network size, and hardware platform. Nevertheless, Figure 15 evaluates direct extrapolation, full-domain retraining, and transfer learning within the same task, thereby providing an internally consistent comparison of both the recovery of predictive accuracy and the reduction in training cost achieved through transfer learning.

4.6. Model Comparison and Applicability

To determine whether the performance improvement in the 80% enlargement target domain arises from increased network complexity or from the transfer-learning mechanism itself, Table 6 compares the predictive performance of different models under the same high-enlargement condition. For a consistent comparison, all models were evaluated using the same 80% enlargement configuration, identical independent CFD test points, and the same error definitions. The same target-domain CFD data partition was used for model training or adaptation, while TLP-PINNs additionally employed source-domain pretraining. The respective model architectures and optimization strategies were retained. TLP-PINNs achieve the best overall performance, with a relative RMSE of 0.0950 and an R_p^2 of 0.9662 for velocity magnitude, together with a relative RMSE of 0.0134 and an R_p^2 of 0.9992 for pressure. PIGNN still yields positive R_p^2, but its prediction errors are higher than those of TLP-PINNs. By contrast, SA-PINNs and SK-PINNs produce a low R^2 for velocity magnitude, indicating poorer predictive performance.

Table 6. Predictive performance of different models for the 80% enlargement case.

Model	Velocity Magnitude RMSE	Velocity Magnitude R_p^2	Pressure RMSE	Pressure R_p^2
Parametric PINN (direct extrapolation)	0.1350	-	0.0350	-
TLP-PINNs	0.0950	0.9662	0.0134	0.9992
PIGNN	0.1691	0.9419	0.2942	0.9834
SA-PINNs	0.6811	0.8926	0.1336	0.9172
SK-PINNs	0.7449	0.7123	0.4945	0.7811

Geometric transfer-learning studies by Takao [34] similarly showed that effective cross-geometry prediction depends not merely on increased network complexity, but on reusing source-domain representations and selectively adapting the model to new geometries. The present results are consistent with this conclusion. The proposed TLP-PINNs framework combines explicit enlargement-ratio parameterization, partial-layer freezing, sparse target-domain supervision, and extrapolation-distance-driven physics-point allocation. These mechanisms enable the optimization to focus on the local flow redistribution induced by large borehole enlargement. Taken together, Table 6 and Figures 11–13 show that greater network complexity or lower training loss does not necessarily guarantee better out-of-domain generalization. Instead, continuous geometric representation, physics-based constraints, and target-domain transfer learning jointly contribute to improved prediction under high-enlargement conditions.

From a computational perspective, CFD remains the high-fidelity reference method in this study. However, each new wellbore geometry generally requires repeated geometry processing, mesh generation, and iterative numerical solution. Once source-domain training is completed, flow fields within the training range can be predicted directly. For high-enlargement conditions outside the source domain, TLP-PINNs require only regeneration of physics collocation points for the target geometry and updating of a subset of network parameters rather than complete retraining from random initialization.

The computational times reported in Figure 15 show that this strategy substantially reduces the cost of neural-network retraining. However, these results should not be interpreted as a universal end-to-end acceleration factor for the complete CFD-to-PINN workflow. Therefore, TLP-PINNs can reduce repeated neural-network training costs across different wellbore geometries while maintaining accuracy comparable to full-domain retraining. Its overall engineering acceleration relative to the complete CFD workflow requires further quantitative assessment on a unified computational platform.

5. Conclusions

This study developed TLP-PINNs for cross-geometry prediction of three-dimensional non-Newtonian flow fields in locally enlarged irregular wellbores. By integrating geometric parameterization, physics-informed learning, adaptive collocation-point allocation, and transfer learning, the framework enables knowledge learned from low-enlargement cases to be transferred to unseen high-enlargement geometries. The main conclusions are as follows:

(1) The parameterized source model can continuously represent geometry-dependent flow fields within the training range. By introducing the borehole enlargement ratio as an explicit input, a single model trained on the 10%, 30%, and 50% cases provides good interpolation accuracy for the velocity field at intermediate geometries, while the pressure prediction exhibits greater case-to-case variability. Multi-seed tests further confirm stable physics-constrained optimization.

(2) Prediction accuracy deteriorates progressively when extrapolating beyond the source domain. For the 100% enlargement case, the relative RMSE of velocity magnitude

increases to 0.1878 and R_p^2 decreases to 0.8805, demonstrating that strong interpolation performance does not guarantee reliable out-of-domain prediction.

(3) Transfer learning substantially improves prediction accuracy for high-enlargement geometries. For the unseen 80% case, the relative RMSEs of axial velocity, velocity magnitude, and pressure are 9.42%, 9.50%, and 1.34%, respectively. Transfer learning mainly corrects errors near the enlargement center and transition zones, recovering the low-velocity region and pressure redistribution with accuracy close to full retraining.

(4) The proposed strategy also reduces repeated neural-network training costs. Full retraining requires 6340 min, whereas source pretraining followed by transfer requires 3095 min, reducing the total training time by 51.2%. With a pretrained source model available, adaptation to the 80% geometry requires only 368 min. These savings refer specifically to neural-network retraining rather than an overall CFD-to-PINN acceleration factor.

The present study is limited to steady, single-phase Herschel–Bulkley flow in a concentric Gaussian-enlarged wellbore. The predicted low-velocity regions indicate potentially unfavorable flow conditions for hole cleaning, but do not directly demonstrate cuttings deposition, transport efficiency, or removal performance. Further particle–fluid coupling simulations and experimental validation are required to establish the relationship between the predicted flow fields and actual hole-cleaning effectiveness. Future work will consider drillstring eccentricity and rotation, broader geometric and rheological parameters, particle–fluid interactions, and experimental validation, with the aim of extending the framework to cuttings transport and hole-cleaning prediction.

Author Contributions: Methodology, Z.X.; software, J.Y.; validation, C.Z.; formal analysis, Y.Z.; data curation, L.C.; writing—original draft preparation, J.W.; writing—review and editing, C.Z.; visualization, Z.X.; supervision, Z.X.; funding acquisition, C.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was financially supported by the National Natural Science Foundation of China (Grant No. 52374003), the Department of Science and Technology of Hubei Province (Grant No. 2023BCB111), and the Hubei Provincial Department of Education (Grant No. T2021004).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The author would like to thank the editor and the anonymous referees for their helpful comments and suggestions.

Conflicts of Interest: Authors Zhichao Xie, Jian Yang, Long Chen, Jiali Wang and Yunhai Zhao were employed by the company PetroChina Southwest Oil & Gasfield Company. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Akhshik, S.; Behzad, M.; Rajabi, M. CFD–DEM approach to investigate the effect of drill pipe rotation on cuttings transport behavior. *J. Pet. Sci. Eng.* **2015**, *127*, 229–244. [CrossRef]
2. Gulraiz, S.; Gray, K.E. Investigating the effects of plug viscosity on annular pressure drop and cuttings transport in a concentric annulus. *J. Nat. Gas Sci. Eng.* **2020**, *76*, 103210. [CrossRef]
3. Naccache, M.F.; Pinto, H.A.M.; Abdu, A.A.Q. Flow displacement in eroded regions inside annular ducts. *J. Braz. Soc. Mech. Sci. Eng.* **2018**, *40*, 420. [CrossRef]

4. Renteria, A.; Maleki, A.; Frigaard, I.A.; Lund, B.; Khadrbeik, M.A.T.; Ytrehus, J.D.M. Effects of irregularity on displacement flows in primary cementing of highly deviated wells. *J. Pet. Sci. Eng.* **2019**, *172*, 662–680. [CrossRef]
5. Skadsem, H.J.; Kragset, S.; Lund, B.; Ytrehus, J.D.; Taghipour, A. Annular displacement in a highly inclined irregular wellbore: Experimental and three-dimensional numerical simulations. *J. Pet. Sci. Eng.* **2019**, *172*, 998–1013. [CrossRef]
6. Etrati, A.; Roustaei, A.; Frigaard, I.A. Strategies for mud-removal from washouts during cementing of vertical surface casing. *J. Pet. Sci. Eng.* **2020**, *195*, 107454. [CrossRef]
7. Renteria, A.; Sarmadi, P.; Thompson, C.; Frigaard, I.A. Effects of wellbore irregularity on primary cementing of horizontal wells. Part 1: Large scale effects. *J. Pet. Sci. Eng.* **2022**, *208*, 109581. [CrossRef]
8. Varges, P.R.; Rodrigues, E.C.; Moraes, L.; de Souza Mendes, P.R.; Naccache, M.F. Flow instabilities in fluid displacement through enlarged regions in annular ducts. *J. Non-Newton. Fluid Mech.* **2022**, *305*, 104834. [CrossRef]
9. 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]
10. 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]
11. Zhao, C.; Zhang, F.; Lou, W.; Wang, X.; Yang, J. A comprehensive review of advances in physics-informed neural networks and their applications in complex fluid dynamics. *Phys. Fluids* **2024**, *36*, 101301. [CrossRef]
12. Pan, C.; Li, C.; Liu, Y.; Zheng, Y.; Ye, H. SK-PINN: Accelerated physics-informed deep learning by smoothing kernel gradients. *Comput. Methods Appl. Mech. Eng.* **2025**, *440*, 117956. [CrossRef]
13. Chen, H.; Xue, L.; Liu, L.; Zou, G.; Han, J.; Dong, Y.; Cong, M.; Liu, Y.; Hosseini-Nasab, S.M. Physics-informed graph neural network for predicting fluid flow in porous media. *Pet. Sci.* **2025**, *22*, 4240–4253. [CrossRef]
14. Zhao, C.; Zhang, F.; Ren, Y.; Lou, W.; Chen, J. Enhanced self-adaptive physics-informed neural networks for prediction of wellbore fluid thermal dynamics during deepwater drilling. *Ocean Eng.* **2026**, *365*, 127161. [CrossRef]
15. Gao, H.; Sun, L.; Wang, J.-X. PhyGeoNet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state PDEs on irregular domain. *J. Comput. Phys.* **2021**, *428*, 110079. [CrossRef]
16. Peng, J.-Z.; Hua, Y.; Aubry, N.; Chen, Z.-H.; Mei, M.; Wu, W.-T. Physics-informed graph convolutional neural network for modeling geometry-adaptive steady-state natural convection. *Int. J. Heat Mass Transf.* **2023**, *216*, 124593. [CrossRef]
17. Xu, C.; Cao, B.T.; Yuan, Y.; Meschke, G. Transfer learning based physics-informed neural networks for solving inverse problems in engineering structures under different loading scenarios. *Comput. Methods Appl. Mech. Eng.* **2023**, *405*, 115852. [CrossRef]
18. Liu, K.; Luo, K.; Cheng, Y.; Liu, A.; Li, H.; Fan, J.; Balachandar, S. Parameterized physics-informed neural networks (P-PINNs) solution of uniform flow over an arbitrarily spinning spherical particle. *Int. J. Multiph. Flow* **2024**, *180*, 104937. [CrossRef]
19. Yan, M.; Liu, K.; An, S.; Wang, G.; Jin, Z.; Zang, J. Transfer learning-driven adaptive physics-informed neural networks acceleration method for aerodynamic identification. *Eng. Appl. Artif. Intell.* **2026**, *166*, 113589. [CrossRef]
20. Wang, L.; Dong, M.; Wang, L.; Huang, C.; Song, D.; Fan, X.; Yang, J.; Wang, T.; Chen, S.; Li, Q. Multi-scale wake modeling based on physics-informed neural networks and transfer learning. *Appl. Energy* **2026**, *406*, 127318. [CrossRef]
21. Jagtap, A.D.; Kawaguchi, K.; Karniadakis, G.E. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *J. Comput. Phys.* **2020**, *404*, 109136. [CrossRef]
22. Wang, S.; Teng, Y.; Perdikaris, P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM J. Sci. Comput.* **2021**, *43*, A3055–A3081. [CrossRef]
23. Wang, S.; Yu, X.; Perdikaris, P. When and why PINNs fail to train: A neural tangent kernel perspective. *J. Comput. Phys.* **2022**, *449*, 110768. [CrossRef]
24. McClenny, L.D.; Braga-Neto, U.M. Self-adaptive physics-informed neural networks. *J. Comput. Phys.* **2023**, *474*, 111722. [CrossRef]
25. 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. *Comput. Methods Appl. Mech. Eng.* **2023**, *403*, 115671. [CrossRef]
26. Aliakbari, M.; Mahmoudi, M.; Vadasz, P.; Arzani, A. Predicting high-fidelity multiphysics data from low-fidelity fluid flow and transport solvers using physics-informed neural networks. *Int. J. Heat Fluid Flow* **2022**, *96*, 109002. [CrossRef]
27. Wang, J.; Huang, L.; Liang, Y.; Dong, J.; Sun, W. Transfer learning enhanced physics-informed neural network for buried pipeline deformation analysis under permanent ground deformation. *Comput. Geotech.* **2026**, *189*, 107630. [CrossRef]
28. Qu, J.; Zhu, H.; Sun, X.; Yan, T. Influence of drilling fluid rheology on hole cleaning: A comprehensive review. *Geoenergy Sci. Eng.* **2025**, *254*, 214030. [CrossRef]
29. Bui, B.T.; Luong, L.H. Effect of field viscometer dial readings on the estimation of drilling fluid frictional pressure loss using Herschel–Bulkley model. *Geoenergy Sci. Eng.* **2024**, *236*, 212697. [CrossRef]
30. Weller, H.G.; Tabor, G.; Jasak, H.; Fureby, C. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.* **1998**, *12*, 620–631. [CrossRef]
31. Luo, J.; Yang, Y.; Yuan, Y.; Xu, S.; Hao, W. An imbalanced learning-based sampling method for physics-informed neural networks. *J. Comput. Phys.* **2025**, *534*, 114010. [CrossRef]

32. Chen, W.; Howard, A.A.; Stinis, P. Self-adaptive weights based on balanced residual decay rate for physics-informed neural networks and deep operator networks. *J. Comput. Phys.* **2025**, *542*, 114226. [CrossRef]
33. Gao, B.; Yao, R.; Li, Y. Physics-informed neural networks with adaptive loss weighting algorithm for solving partial differential equations. *Comput. Math. Appl.* **2025**, *181*, 216–227. [CrossRef]
34. Takao, R.; Ii, S. Fine-tuning physics-informed neural networks for cavity flows using coordinate transformation. *Comput. Fluids* **2026**, *306*, 106957. [CrossRef]
35. Zhu, K.; Chen, X.; Ao, X.; He, Z. Parameter-efficient transfer of physics-informed neural networks for buoyancy-driven enclosures via geometry-conditioned adapters. *Phys. Fluids* **2026**, *38*, 023614. [CrossRef]
36. Liu, L.; Luo, J.; Heng, Y. Fast solution of 3D transport processes using a physics-informed neural network with embedded transfer learning. *Sep. Purif. Technol.* **2025**, *363*, 132012. [CrossRef]
37. Qu, J.; Zhu, H.; Wang, Q.; Sun, X.; Yan, T. A method for determining minimum flow rate for cuttings transport in upper large-diameter section of ultra-deep wells. *Geoenergy Sci. Eng.* **2025**, *258*, 214307. [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.