Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Modeling and simulating turbulent secondary flow in a streamwise-rotating channel with LLM-guided symbolic regression and physics-informed neural networks

Wei Zhang

waynezw0618@163.com

Marine Design and Research Institute of China https://orcid.org/0000-0001-9828-0829

Benlong wang

Shanghai Jiao Tong University School of Naval Architecture Ocean and Civil Engineering

zheng Ma

CSSRC: China Ship Scientific Research Center

Bing Zhu

University of Shanghai for Science and Technology School of Energy and Power Engineering

Research Article

Keywords: secondary flow, Taylor–Görtler vortex, symbolic regression, large language model, algebraic stress model, physics-informed neural network

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11375175/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Springer

Available online at https://link.springer.com/journal/42241
http://www.jhydrodynamics.com
Journal of Hydrodynamics, XXX, XXX(XXX): XXX–XXX
https://doi.org/10.1007/s42241-XXX-XXXX-X

Modeling and simulating turbulent secondary flow in a streamwise-rotating channel with LLM-guided symbolic regression and physics-informed neural networks

Wei Zhang[1*], Benlong Wang[2], Zheng Ma[3], Bing Zhu[4]

1. Marine Design and Research Institute of China, Shanghai, 200011, China
2. School of Naval Architecture, Ocean and Civil Engineering, Shanghai Jiao Tong University, Shanghai, 200240, China
3. China Ship Scientific Research Center, Shanghai, 200011, China
4. School of Energy and Power Engineering, University of Shanghai for Science and Technology, Shanghai, 200093, China

(Received month, date, year, Revised month, date, year, Accepted month, date, year, Published online month, date, year)
© China Ship Scientific Research Center 2026

Abstract: Turbulent secondary flow in a streamwise-rotating channel is dominated by Taylor–Görtler (TG) vortices whose direct numerical simulation is expensive. This paper develops a two-stage artificial-intelligence framework that *learns the closure* and *simulates the flow*. In the first stage, the LGF (LLM-guided Graph-Feynman) equation-discovery engine targets the single decisive anisotropy component, the secondary-flow covariance $b_{23} = \langle vw \rangle / 2k$: twenty discovery runs (four feature pools of rotation-variant tensor channels, two LGF operating arms, and PySR at three seeds) are scored against DNS statistics of five rotation numbers $Ro_\tau = 7.5$–150. The guided search buys stability: the eight LGF runs hold the case-normalized error within 0.051–0.137 at 13–17 nodes, whereas the twelve PySR runs spread over 0.042–0.334 with seed-level variation up to eightfold at about 60% more nodes. The four selected closures (0.042–0.055, five to seven constants each) share one skeleton: a saturating function of the local rotation number, balancing Coriolis production against isotropization, crossed with a wall function of the wall distance. In the second stage they are carried through the mean-spanwise-profile balance into a physics-informed spectral network (PINN) that reconstructs the time-averaged TG-vortex field at $Ro_\tau = 7.5$, 30 and 75 in minutes on a desktop CPU. All four reproduce the two-layer counter-rotating roll cells with secondary-flow amplitudes within 10% of the DNS despite mean-profile errors of order the profile itself: the reconstruction is robust to O(1) profile errors, identifying stress production and viscosity, not mean-flow advection, as the governing balance and completing a "discover–close–simulate" pipeline with compact learned closures.

Key words: secondary flow, Taylor–Görtler vortex, symbolic regression, large language model, algebraic stress model, physics-informed neural network

0. Introduction

Turbulent secondary flows (weak, statistically steady circulations superimposed on the primary shear) control momentum, heat and scalar transport in ducts, turbomachinery passages and marine-propulsion components, and are therefore of long-standing interest to the hydrodynamics community[1]. A canonical configuration in which they appear in isolation is the plane channel rotating about the streamwise (x) axis, where the Coriolis force drives a mean spanwise flow $\bar{W}(y)$ whose interaction with the turbulence organizes streamwise-elongated, counter-rotating roll cells known as Taylor–Görtler (TG) vortices[2-3]. Unlike spanwise rotation, which produces a single layer of roll cells[2], streamwise rotation arranges the TG vortices into *two* layers of counter-rotating pairs stacked across the channel[3], a signature that makes the configuration a stringent test bed for both modeling and simulation.

The physical understanding of these vortices rests almost entirely on DNS: linear and weakly nonlinear analyses established the instability mechanism[4-6], DNS and group analysis documented the mean-flow reversal and roll-cell structure[7-10], and Yang and Wang[3] quantified the vortex scales at $Re_\tau = 180$,

Project supported by the National Natural Science Foundation of China (Grant Nos. U2341242, 12272357).
Biography: Wei Zhang (1981–), Male, Ph.D.,
E-mail: waynezw0618@163.com
Corresponding author: Wei Zhang,
E-mail: waynezw0618@163.com

Ro_τ = 7.5–150. The spanwise wavelength is nearly rotation-independent, but the streamwise extent grows so rapidly that $Ro_\tau = 150$ requires $L_x = 512\pi h$, so systematic DNS is expensive and low-cost surrogates valuable.

Two complementary strands of artificial intelligence (AI) research are relevant. The first is *equation discovery*: symbolic regression and its physics-informed and, more recently, large-language-model (LLM) driven variants distill interpretable constitutive laws and stress closures from data[11-18]. For turbulence, Ling et al.[19] embedded Galilean invariance into a neural Reynolds-stress closure, and Weatheritt and Sandberg[20] discovered an ASM for rotating flow with an evolutionary algorithm acting on the Pope tensor basis[21-22]; the data-driven construction and calibration of such closures is itself an active field[23-24]. The second strand is *equation solving*: physics-informed neural networks (PINNs)[25-26] embed governing equations in the training loss and handle forward and inverse fluid problems, including data assimilation and turbulence modeling[27-29]. Their well-documented training difficulties[30-31] remain a practical concern.

The present work connects the two strands into a single "discover–close–simulate" pipeline and applies it to the streamwise-rotating channel; our contributions are as follows.

1. We develop LGF (LLM-guided Graph-Feynman), a hybrid equation-discovery engine that represents candidate equations as edge-featured directed acyclic graphs, injects AI-Feynman-style physics priors (dimensional, symmetry, separability and compositionality tests), and couples an outer LLM "semantic" loop with an inner graph-genetic-programming loop for parameter fitting and selection.
2. Using LGF we infer the closure for the single decisive anisotropy component, $b_{23} = \langle vw\rangle/2k$, from the DNS statistics of Yang and Wang[3], over eight runs spanning four feature pools and two operating arms, and benchmark it against twelve runs of pure symbolic regression (PySR).
3. We carry the discovered closures through the mean-spanwise-profile balance into a physics-informed spectral network that reconstructs the time-averaged TG-vortex field in the (y, z) plane from one-dimensional profiles with a hard-constraint spectral ansatz and a scale-invariant loss, and we quantify how the closure error propagates end to end into the simulated secondary flow.

1. Physical problem and data

1.1 Flow configuration and governing equations

We consider turbulent channel flow between two infinite parallel walls at $y = \pm h$, driven by a mean pressure gradient in x and rotating about the streamwise x-axis with angular velocity Ω. The friction Reynolds and rotation numbers are $Re_\tau = u_\tau h/\nu = 180$ and $Ro_\tau = 2\Omega h/u_\tau$, where u_τ is the friction velocity, ν the kinematic viscosity and h the half-channel height. In the fully developed state all statistics depend on y only, the secondary flow additionally on z; hence $\partial/\partial x = 0$. All quantities are nondimensionalized by h and u_τ.

The mean streamwise vorticity $\Omega_x = \partial\bar{W}/\partial y - \partial\bar{V}/\partial z$ obeys the equation given by Bradshaw[1]. Under fully developed conditions the vortex-stretching and skew-induced terms vanish, and the Coriolis force contributes nothing to the x-component (the curl of solid-body rotation is divergence-free); the rotation acts indirectly, through the strongly rotation-dependent DNS profiles. Writing the secondary flow with a stream function $\phi(y, z)$,

$$\tilde{v} = \frac{\partial\phi}{\partial z}, \quad \tilde{w} = -\frac{\partial\phi}{\partial y}, \quad V = \tilde{v}, \quad W = \bar{W}(y) + \tilde{w}, \quad (1)$$

so that $\Omega_x = \mathrm{d}\bar{W}/\mathrm{d}y - \nabla^2_{yz}\phi$ and incompressibility holds identically, the surviving mode-2 balance is

$$\begin{aligned}&\left[V\frac{\partial\Omega_x}{\partial y} + W\frac{\partial\Omega_x}{\partial z}\right] - \left\langle V\frac{\partial\Omega_x}{\partial y} + W\frac{\partial\Omega_x}{\partial z}\right\rangle_z \\ &= \left(\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2}\right)\left(-\langle vw\rangle_{\mathrm{vtx}}\right) \\ &\quad + \frac{\partial^2\langle v^2 - w^2\rangle_{\mathrm{vtx}}}{\partial y\,\partial z} + \nu\nabla^2_{yz}\tilde{\Omega}_x,\end{aligned} \quad (2)$$

where $\tilde{\Omega}_x = -\nabla^2_{yz}\phi$; the subscript "vtx" denotes the z-dependent part of the Reynolds stresses generated by the roll cells, while the z-independent part is absorbed by the mode-1 balance $\nu\,\mathrm{d}^2\bar{W}/\mathrm{d}y^2 = \mathrm{d}\langle vw\rangle/\mathrm{d}y$ with $\bar{W}(\pm1) = 0$, which determines the mean spanwise profile $\bar{W}(y)$ once $\langle vw\rangle(y)$ is specified.

1.2 The closure problem

Equation (2) is not closed: the vortex self-stress $\langle v_i v_j\rangle_{\mathrm{vtx}}$ must be related to the resolved motion, and the mean profile $\bar{W}(y)$ must be supplied. Both requirements lead to the same anisotropy component. We seek an explicit closure for the Reynolds-stress anisotropy tensor

$$b_{ij} = \frac{\langle u_i u_j\rangle}{2k} - \frac{1}{3}\delta_{ij}, \quad (3)$$

with $k = \langle u_i u_i\rangle/2$ and $\mathrm{tr}(b) = 0$. Following Pope[21,32], b_{ij} is expanded on the ten symmetric traceless tensor

bases $T_{ij}^{(m)}$ constructed from the normalized mean strain $\boldsymbol{S}$ and rotation $\boldsymbol{W}$; explicit algebraic closures of this type trace back to the Reynolds-stress closure of Launder et al.[33] and were brought to a self-consistent explicit form by Wallin and Johansson[34]:

$$b_{ij} = \sum_{m=1}^{10} g_m(I_1, \ldots, I_5, Re_t, Ro_{\mathrm{loc}})\, T_{ij}^{(m)}, \tag{4}$$

where the five invariants are $\eta^2 = \mathrm{tr}(\boldsymbol{S}^2)$, $\zeta^2 = \mathrm{tr}(\boldsymbol{W}^2)$, $s_3 = \mathrm{tr}(\boldsymbol{S}^3)$, $so_2 = \mathrm{tr}(\boldsymbol{S}^2\boldsymbol{W})$, $s_2o_2 = \mathrm{tr}(\boldsymbol{S}^2\boldsymbol{W}^2)$, together with $Re_t = k^2/(\nu\varepsilon)$ and $Ro_{\mathrm{loc}} = 2\Omega k/\varepsilon$.

Of the six independent components of b_{ij}, one is decisive for the present problem: the secondary-flow covariance $b_{23} = \langle vw\rangle/2k$. Its flux closes the mode-1 balance of Sec. 2.1 (verified against the DNS to within 0.4% rms on $\bar{W}$), which fixes the mean spanwise flow $\bar{W}(y)$, and through $\bar{W}$ and its wall gradient it enters the advection and the base vorticity of the mode-2 balance, Eq. (2). Rather than inferring all ten coefficient functions g_m at once (the full ASM of Weatheritt and Sandberg[20]), we therefore pose the discovery problem on the single decisive component, $b_{23} = f(I_1, \ldots, I_5, Re_t, Ro_{\mathrm{loc}}, y_w, \{T_{23}^{(m)}\})$, with $y_w = 1 - |y|/h$ the dimensionless wall distance appended to the Pope feature space and $\{T_{23}^{(m)}\}$ ($m = 1, 5, 6, 10$) the four symmetry-allowed tensor channels through which the bases of Eq. (4) can feed the 23 component. The function f is then learned directly from the DNS statistics of the five rotation numbers.

1.3 DNS database

The one-dimensional DNS statistics of Yang and Wang[3] at $Re_\tau = 180$ are used throughout: five rotation numbers $Ro_\tau = 7.5$–150 provide $U^+(y)$, $W^+(y)$ and the full $\langle u_i u_j\rangle^+(y)$ on 129 wall-normal nodes. These profiles feed both AI stages: LGF learns the g_m across the five cases; the PINN uses the $Ro_\tau = 7.5$, 30 and 75 profiles to reconstruct the TG-vortex field.

2. AI methodology

2.1 LGF: LLM-guided Graph-Feynman equation discovery

LGF is a hybrid engine that unifies three equation-discovery paradigms in a hierarchical double loop (Fig. 1). Its building blocks are:

- *Graph representation.* Every candidate equation is a directed acyclic graph whose nodes are operators or variables and whose edges carry features (constants or per-case free parameters), so that free coefficients and topology live in one representation, a capability string- or tree-based search lacks[18].
- *Physics priors.* AI-Feynman-style tests (dimensional analysis, symmetry, separability, compositionality) are converted into graph constraints that prune the search space before evolution[15].
- *LLM semantic operator.* An LLM proposes or mutates equation skeletons at the semantic level (e.g. injecting a tanh saturation to suppress unbounded growth), guided by error attribution that localizes the offending subgraph[17]; using an LLM to drive the search over program space follows the program-search paradigm of FunSearch[35], and elite graphs and scores are reinjected into the prompt through an experience buffer.
- *Graph genetic programming.* An inner fast loop performs subgraph crossover, node/feature mutation, extinction and top-k selection, with L-BFGS multi-start fitting of the edge parameters of every candidate graph.

The LLM is a GLM-5.3 model with two modes: *G_INIT* generates the initial population of equation skeletons, and *M_GUIDE* proposes targeted mutations during evolution. Two operating arms are compared: "go" (G_INIT only, then pure graph-GP evolution) and "mg" (G_INIT plus M_GUIDE throughout). All prompt templates, random seeds and per-run statistics are recorded and will be released with the code.

For comparison, the same dataset and features are fed to PySR[16], a pure symbolic-regression engine with no LLM guidance or physics priors, which isolates the contribution of the semantic and prior layers. Both engines search the same target, the closure $b_{23} = f(I_1, \ldots, I_5, Re_t, Ro_{\mathrm{loc}}, y_w, \{T_{23}^{(m)}\})$, on the same DNS database. Four feature pools are explored, built on the tensor bases and invariants of the three rotation-variant mean-rotation tensors $\boldsymbol{W} = \overline{W}$ (v, mean rotation rate), $\boldsymbol{W} = \overline{W} + \boldsymbol{\Omega}$ (g, Gatski–Speziale) and $\boldsymbol{W} = \overline{W} + 2\boldsymbol{\Omega}$ (e, absolute vorticity)[22], plus a mixed pool; $\boldsymbol{S}$ is common to all. Every pool contains the scalar features $\{\eta^2, \zeta^2, s_3, so_2, s_2o_2, Re_t, Ro_{\mathrm{loc}}, y_w\}$ and the four symmetry-allowed channels $T_{23}^{(m)}$ ($m = 1, 5, 6, 10$) taken from the pool's bases; $T_{23}^{(1)} = S_{23} = \frac{1}{2}\,\mathrm{d}\bar{W}/\mathrm{d}y$ is the mean spanwise shear, the only linear production channel of $\langle vw\rangle$ and the channel a linear eddy-viscosity closure would retain ($b_{23} \propto S_{23}$). The v, g and e pools use the respective variant's channels (12 features each); the mixed pool concatenates the three variants (54 features). Each pool is run with LGF in both operating arms and with PySR at three random seeds, giving twenty candidate runs. Candidates are scored by the case-normalized loss

$$\mathcal{L}_{\mathrm{asm}} = \left\langle \frac{(\hat{b}_{23} - b_{23}^{\mathrm{DNS}})^2}{\mathrm{var}_c} \right\rangle, \tag{5}$$

where the hat denotes the candidate model, var_c is the variance of the DNS b_{23} over the wall-normal nodes of

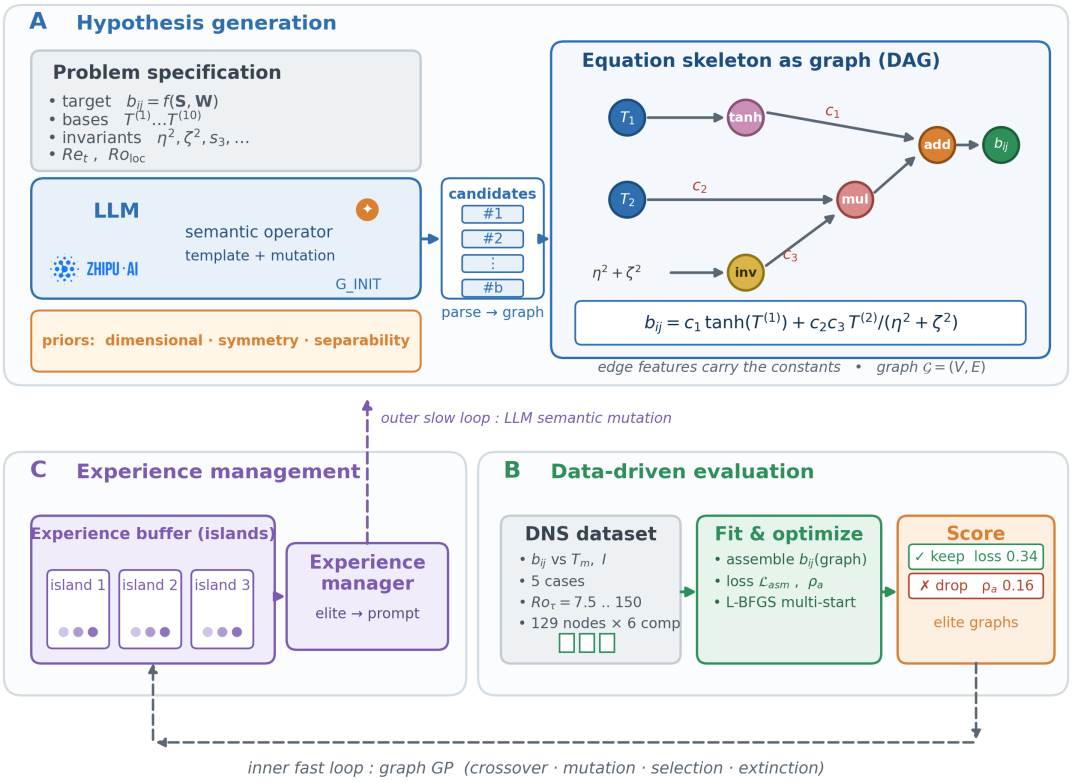


Fig. 1 The LGF (LLM-guided Graph-Feynman) equation-discovery framework: edge-featured DAG candidates, physics priors, an outer LLM semantic loop with an inner graph-GP loop, and an experience buffer feeding elite graphs back to the LLM.

case c, and the average runs over the five cases; and by the shape correlation

$$\rho = \left\langle \mathrm{corr}_y\left(\hat{b}_{23},\, b_{23}^{\mathrm{DNS}}\right) \right\rangle, \tag{6}$$

the per-case Pearson coefficient along y. Because b_{23} is a bias-dominated signal (its mean-to-standard-deviation ratio is 1.6–2.9), a high ρ certifies only the profile shape; ranking by $\mathcal{L}_{\mathrm{asm}}$ is what penalizes the amplitude and offset errors that ρ cannot see, so the two metrics are reported together.

The four best candidates overall are carried through the rest of the paper: three LGF closures from the go arm (the e-pool *exp* and g-pool *tanh* models, mean loss 0.055 each, and the mixed-pool *rat* model, 0.051) and the best PySR model (mixed pool, 0.042). The two engines split the pools (LGF wins g and e, the best PySR seed wins v and m), and the guided search buys stability rather than the single best score: all eight LGF runs land within 0.051–0.137 at 13–17 nodes, against twelve PySR runs spreading over 0.042–0.334 at 12–30 nodes. The go arm also beats the full-guidance mg arm in every pool, so the LLM's decisive contribution here is the physics-informed initialization rather than mid-run mutation. The four selected closures are shown in Eqs. (7)–(10):

$$b_{23}^{(\mathrm{exp})} = c_1\, e^{c_2 y_w} \tanh\left(c_3 Ro_{\mathrm{loc}}\right) \times \left[\ln\left(y_w + c_4\right) - c_5 Ro_{\mathrm{loc}}\right], \tag{7}$$

$$b_{23}^{(\mathrm{tanh})} = c_1 \tanh\left(c_2 y_w + c_4\right) \times \sqrt{Ro_{\mathrm{loc}}} - c_3 Ro_{\mathrm{loc}} - c_5, \tag{8}$$

$$b_{23}^{(\mathrm{rat})} = c_1 \tanh\left(c_2 Ro_{\mathrm{loc}}\right) \times \frac{c_3 y_w + c_4}{c_5 Ro_{\mathrm{loc}} + c_6 y_w}, \tag{9}$$

$$b_{23}^{(\mathrm{PySR})} = \sqrt{Ro_{\mathrm{loc}}}\left(c_1 + c_2 y_w - c_3 y_w^2\right) - c_4\left(c_5 T_{1s}^2 + c_6 T_{1s} + Ro_{\mathrm{loc}} - c_7\right), \tag{10}$$

with the constants of Table 1 (multi-start L-BFGS refits, globally fitted across the five rotation numbers) and $T_{1s} = S_{23} = \frac{1}{2}\,\mathrm{d}\bar{W}/\mathrm{d}y$ the normalized mean spanwise shear. Every closure factorizes into a saturating function of Ro_{loc} crossed with a wall function of y_w alone. The saturation expresses the balance between Coriolis production of $\langle vw \rangle$ and isotropization/dissipation, an intermediate between low-Ro linear growth and high-Ro saturation; it is realized either as the $\sqrt{Ro_{\mathrm{loc}}}$ leading term (the *tanh* closure and the PySR model) or as a $\tanh(cRo_{\mathrm{loc}})$ factor (the *exp* and *rat* closures). One of

Table 1 Fitted constants of the four discovered b_{23} closures (Eqs. (7)–(10)), globally fitted across the five rotation numbers by multi-start L-BFGS.

Model	Constants
exp	c_1=1.72×10^{-1}, c_2=−7.91×10^{-1}, c_3=1.00×10^{-1}, c_4=1.158, c_5=1.94×10^{-3}
tanh	c_1=1.74×10^{-2}, c_2=1.399, c_3=1.08×10^{-3}, c_4=5.70×10^{-1}, c_5=2.24×10^{-4}
rat	c_1=1.074, c_2=7.10×10^{-2}, c_3=2.49×10^{-1}, c_4=2.28×10^{-3}, c_5=2.14×10^{-2}, c_6=3.71
PySR	c_1=8.96×10^{-3}, c_2=1.40×10^{-2}, c_3=6.67×10^{-3}, c_4=1.04×10^{-3}, c_5=−1.79×10^{-2}, c_6=−1.99, c_7=5.82×10^{-3}

the two forms recurs in every one of the eight LGF runs, evidence that the search converges to the physics rather than to a lucky formula. The crossed wall functions provide the near-wall damping that enforces the no-slip tendency of b_{23}: the S-shaped $\tanh(c_2 y_w + c_4)$, the exponential–logarithmic product $e^{c_2 y_w}\ln(y_w + c_4)$, the rational $(c_3 y_w + c_4)/(c_5 Ro_{\rm loc} + c_6 y_w)$ and the quadratic $c_1 + c_2 y_w - c_3 y_w^2$; all are even functions of y by construction, as the true profile is. The PySR model adds a small correction ($\approx 4\%$ of the signal rms) in the mean-shear channel T_{1s}: guided and unguided engines thus agree on the leading balance, and the unguided one independently recovers the linear production channel as a perturbation. All four closures hold five to seven constants against the ten coefficient functions of the full ASM.

2.2 Embedding the b_{23} closures in a physics-informed network

The two AI stages are linked by the mean spanwise profile. Stage one returns an explicit b_{23} closure; evaluated on the DNS mean state it implies a stress profile $\langle vw\rangle_{\rm ASM}(y) = 2k\, b_{23}$, which is carried through the mode-1 balance $\nu\, \mathrm{d}^2\bar{W}/\mathrm{d}y^2 = \mathrm{d}\langle vw\rangle/\mathrm{d}y$, $\bar{W}(\pm 1) = 0$, to yield the modelled mean profile $\bar{W}_{\rm ASM}(y)$. Stage two solves the mode-2 balance of Eq. (2) with a hard-constraint spectral network, taking $\bar{W}_{\rm ASM}$ as the base flow, while the eddy viscosity $\nu_t(y)$ and the energy-calibration target $\langle vv\rangle$ come from the DNS. The closure thus acts on the simulated field *only* through the mean profile it implies; no coefficient is passed to the network directly, and the two stages are strictly decoupled.

This routing makes the error propagation explicit and measurable. The mode-1 integration is exact when fed the DNS $\langle vw\rangle$ (0.4% rms on $\bar{W}$), so any error in $\bar{W}_{\rm ASM}$ is attributable to the closure alone; and $\bar{W}$ enters the mode-2 balance only through the advection $W\,\partial\Omega_x/\partial z$ and the base vorticity $\mathrm{d}\bar{W}/\mathrm{d}y$. The vortex amplitude itself is set by the energy calibration, as follows from the degeneracy of Eq. (2): the equation is homogeneous in ϕ, so the PDE residual constrains the vortex shape but not its amplitude, which is instead fixed by the RMS energy loss against $\langle vv\rangle$. Section 2.3 details the network and the loss.

2.3 PINN: a physics-informed spectral network for the secondary flow

The vortex field $\phi(y,z)$ is obtained by minimizing the residual of Eq. (2) with a PINN. The vortex self-stress is closed by a Boussinesq-type eddy-viscosity model driven by the resolved fluctuation shear,

$$-\langle vw\rangle_{\rm vtx} = \nu_t\left(\phi_{zz} - \phi_{yy}\right), \qquad \langle v^2 - w^2\rangle_{\rm vtx} = -4\,\nu_t\,\phi_{yz}, \tag{11}$$

where the factor 4 follows from the Boussinesq normal stresses, $\langle v^2 - w^2\rangle_{\rm vtx} = -2\nu_t(S_{22} - S_{33})$ with $S_{22} - S_{33} = 2\phi_{yz}$. Here $\nu_t(y)$ is extracted from the DNS one-dimensional momentum balance, $\nu_t = -\langle uv\rangle/(\mathrm{d}\bar{U}/\mathrm{d}y)$ (capped near the walls), and the mean profile $\bar{W}(y)$ is supplied by the closure through the mode-1 balance of Sec. 2.2. The viscous term retains the molecular ν, so ν_t enters the residual only through Eq. (11); both stresses vanish at $\phi = 0$, so the closure cannot sustain the vortex by itself. The stream function is represented by the hard-constraint spectral ansatz

$$\phi(y,z) = (1-y^2)^2 \sum_{k=1}^{m} \left[a_k(y)\sin kz^* + b_k(y)\cos kz^*\right], \tag{12}$$

with $z^* = 2\pi z/\lambda_3$ and $m = 6$ harmonics, so that spanwise periodicity, no-slip and in-plane incompressibility are satisfied *identically*, and the $k \geq 1$-only spectrum reproduces the DNS mean profiles exactly. Only the $2m$ profiles $\{a_k, b_k\}$ are learned, by a three-layer (width 48, tanh) fully connected network (Fig. 2).

Equation (2) is homogeneous in ϕ, so a naive least-squares residual collapses to the trivial solution. Two loss-design devices cure this. First, the PDE loss is made *scale-invariant*,

$$\mathcal{L}_{\rm pde} = \frac{\langle R^2\rangle}{\langle C^2\rangle + \langle \mathcal{S}_1^2\rangle + \langle \mathcal{S}_2^2\rangle + \langle(\nu\nabla^4\phi)^2\rangle}, \tag{13}$$

where C is the z-fluctuating convection term on the left-hand side of Eq. (2), $\mathcal{S}_1$ and $\mathcal{S}_2$ the two stress-production terms on its right-hand side, and $\nu\nabla^4\phi = -\nu\nabla^2\tilde{\Omega}_x$ the viscous term; the denominator is the mean square of the individual terms, so $\mathcal{L}_{\rm pde}$ rewards only the consistency of the vortex *shape* with the equation. Second, the amplitude, which the homogeneous equation cannot determine, is fixed by an RMS energy-calibration loss against the DNS wall-normal stress,

$$\mathcal{L}_{\rm e} = \frac{1}{U_s^2}\left\langle\left(\sqrt{\langle\tilde{v}^2\rangle_z + \varepsilon} - \sqrt{\beta\,\langle vv\rangle_{\rm DNS}}\right)^2\right\rangle_y, \tag{14}$$

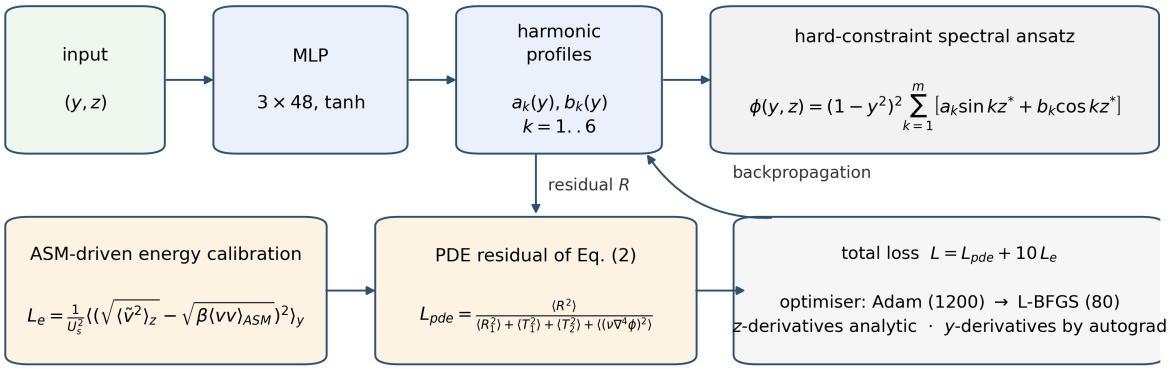


Fig. 2 Physics-informed spectral network for the TG vortices: the MLP maps y to the harmonic profiles $\{a_k, b_k\}$; spectral synthesis and analytic differentiation yield the vortex field and the residual of Eq. (2). The scale-invariant loss shapes the vortex; the ASM-driven energy-calibration loss fixes its amplitude.

with $U_s^2 = \beta \max_y \langle vv \rangle$, $\beta = 0.05$ and the square root guarded near the walls. The total loss $\mathcal{L} = \mathcal{L}_{\text{pde}} + 10\,\mathcal{L}_{\text{e}}$ is minimized by Adam (1200 steps) and L-BFGS (80 iterations); all z-derivatives are evaluated analytically from the spectral basis and all y-derivatives by nested automatic differentiation, and a single run takes 8–15 min on a desktop CPU. The use of a low-order statistical constraint to regularize an otherwise under-determined inverse problem follows the same philosophy as profile-driven PINN data assimilation[27-28].

3. Results and discussion

3.1 Discovered stress closures

Figure 3 shows a representative LGF search history (go arm, g pool): the best-so-far loss declines monotonically with wall-clock time, punctuated by discrete drops where graph-genetic crossover discovers a more favorable skeleton. The same holds in every pool, and the go arm converges both faster and lower than mg: within 10% of its final loss after 22–72 s against 58–149 s, ending at 0.051–0.110 against 0.073–0.137. After the physics-informed initialization, mid-run LLM mutations thus add wall-clock time without improving the outcome; consistently, all three selected LGF closures come from the go arm.

The twenty candidates are compared in Fig. 4(a) in the loss–correlation plane, with the four closures carried to the simulation stage highlighted. Three observations follow. (i) All candidates reach high shape correlations ($\bar{\rho} = 0.96$–0.99): the even, single-peak structure of b_{23} is within reach of every engine, as anticipated from the bias-dominated nature of the signal. (ii) The case-normalized loss separates the engines by stability: the eight LGF runs occupy a narrow band $\mathcal{L}_{\text{asm}} = 0.051$–$0.137$, whereas the twelve PySR runs scatter over 0.042–0.334 with up to a factor-of-eight

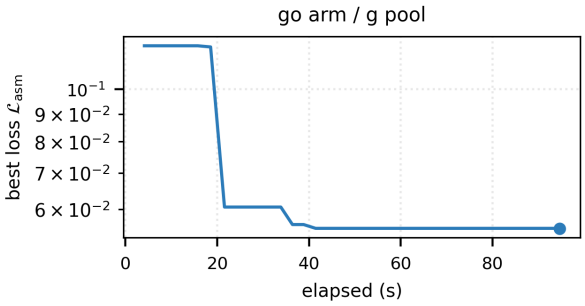


Fig. 3 Representative LGF search history (go arm, g pool): best-so-far loss $\mathcal{L}_{\text{asm}}$ versus wall-clock time (log scale), with discrete drops where graph-genetic crossover discovers a more favorable skeleton.

seed-to-seed spread within one pool; LGF thus wins the g and e pools while the best PySR seed wins the v and m pools. (iii) The four carried forward are the four best candidates overall (0.042–0.055); the *rat* and *exp* closures hold the two highest shape correlations of the campaign ($\bar{\rho} = 0.993$ and 0.992, against 0.991 for PySR). Figure 4(b) and Table 2 break the loss down by rotation number: every selected closure holds its NMSE at or below 0.115 over $Ro_\tau = 7.5$–150, with the best single values of the study (0.004, *rat* at $Ro_\tau = 30$ and PySR at $Ro_\tau = 15$) and the largest residuals at mid-range (0.080–0.115). The highest rotation number is the easiest case for all four (NMSE ≤ 0.031).

3.2 From closure to mean spanwise profile

Figure 5 traces the selected closures through the mode-1 chain at the three rotation numbers retained for the simulation stage. The top row compares the modelled $b_{23}(y)$ with the DNS: all four closures reproduce the even, sign-definite single-peak profile (per-case $\rho = 0.977$–0.999), with the *rat* closure at $Ro_\tau = 30$ and the *exp* closure at $Ro_\tau = 7.5$ tracking the DNS curve almost exactly. The bottom row shows the mean span-

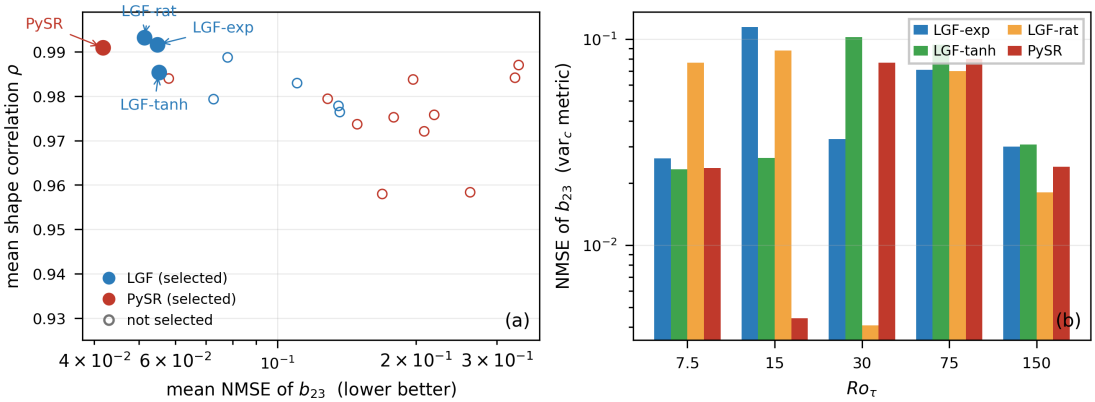


Fig. 4 Assessment of the discovered b_{23} closures. (a) Mean case-normalized loss versus mean shape correlation for the twenty candidate runs (eight LGF, twelve PySR); the four closures carried to the simulation stage are highlighted. (b) Per-rotation-number NMSE of the four selected closures (log scale).

Table 2 Case-normalized loss (NMSE) of the four b_{23} closures at the five rotation numbers, with the mean, the worst case and the mean shape correlation $\bar{\rho}$; bold marks the best value per column.

Model	7.5	15	30	75	150	Mean	Worst	$\bar{\rho}$
exp	0.026	0.115	0.033	0.071	0.030	0.055	0.115	0.9917
tanh	**0.023**	0.026	0.102	0.094	0.031	0.055	0.102	0.9854
rat	0.077	0.088	**0.004**	**0.070**	**0.018**	0.051	0.088	**0.9932**
PySR	0.024	**0.004**	0.077	0.080	0.024	**0.042**	**0.080**	0.9910

wise profiles $\bar{W}(y)$ implied by the closures through the mode-1 balance; fed the DNS $\langle vw \rangle$, this integration reproduces the DNS $\bar{W}$ to 0.2–0.4% rms, so the displayed differences are attributable to the closures alone. The DNS $\bar{W}$ consists of two steep near-wall layers (extrema at $|y| \approx 0.9$, zero crossings at $|y| \approx 0.43$–0.52) around a near-zero core. The *exp* closure is the most stable, its $\bar{W}$ error staying within 0.27–0.35 at all three rotation numbers, and the *rat* closure the most accurate where its b_{23} is best (0.080 at $Ro_\tau = 30$, 0.227 at $Ro_\tau = 7.5$); alone among the four it also reproduces the near-wall layering qualitatively, its interior zero crossings separating the near-wall layers from the core as in the DNS. The tanh and PySR closures degrade at the higher rotation numbers (0.49–1.06 at $Ro_\tau = 30$–75), where their smooth wall functions produce a single broad hump instead of the near-wall layers; the rms error of the modelled $\bar{W}$ thus spans 0.080–1.06 (relative to $\max|\bar{W}|$). The question that matters for the pipeline, however, is how these profile errors propagate into the simulated secondary flow; the answer, given next, is that the propagation is weak.

3.3 Simulating the secondary flow with the AI closures

Finally the four closures drive the PINN of Section 2.3 at $Ro_\tau = 7.5$, 30 and 75; the closure-implied $\bar{W}(y)$ is the only closure-dependent input. Figure 6 compares the DNS baseline with the four reconstructions on a common (y, z) window of one spanwise wavelength, and Table 3 lists the diagnostics. All four closures reproduce the two-layer counter-rotating roll cells at every rotation number; the amplitude $\max|\tilde{v}|$ lands within 1.9–9.4% of the DNS in all twelve runs. At $Ro_\tau = 7.5$ the *exp* closure recovers the DNS layering exactly: four $\tilde{\Omega}_x$ sign changes per wavelength, strongest row at $y/h = -0.25$ against -0.26. The *rat* closure matches the row location exactly at a finer layering, while the tanh and PySR closures select a coarser branch (two sign changes, strongest row in the upper half). At $Ro_\tau = 30$ the *exp* and *rat* closures keep the DNS row position (+0.26 and +0.25 against +0.25) at a slightly coarser layering, and PySR comes closest in the amplitude-type diagnostics ($\max|\tilde{v}|$ and $\max|\tilde{w}|$ within 6%). At $Ro_\tau = 75$, the case with the strongest mean spanwise flow, the tanh closure is the most faithful in structure (strongest row at $y/h = +0.67$ against +0.69 for the DNS, amplitude within 6%), while the *rat* closure is the most faithful in intensity, its vorticity peak erring by less than 2% (9.39 against 9.23). The scale-invariant PDE loss, Eq. (13), spans 0.409–0.578 across the closure-driven runs, bracketing the DNS-baseline values (0.514–0.536), so every reconstruction is an admissible solution of its own closed equation.

The most instructive outcome is the propagation of the profile errors of Sec. 4.2, and the propagation is weak. The smallest and the largest profile errors of the campaign produce comparable amplitude errors: 0.080 (*rat*, $Ro_\tau = 30$) and 1.06 (*tanh*, $Ro_\tau = 75$), a thirteenfold span, give 6.1% and 5.8%. The tanh run with the larger error also places the strongest vorticity row closer to the DNS than any other reconstruction. The amplitude error is thus not monotone in the profile error, and no run exceeds 10%. The roll-cell field is evidently governed by the stress-production and viscous terms of Eq. (2), with the mean-flow advection entering as a correction: the reconstruction inherits its shape from the

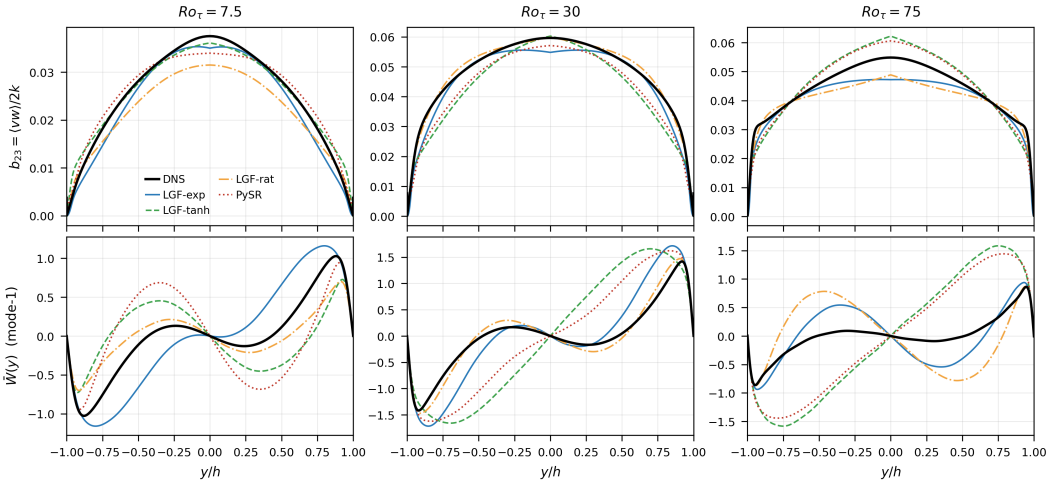


Fig. 5 Closure-induced profiles at $Ro_\tau = 7.5$, 30 and 75 (left to right). Top: modelled $b_{23} = \langle vw \rangle / 2k$ against the DNS. Bottom: mean spanwise profile $\bar{W}(y)$ implied by each closure through the mode-1 balance $\nu\, \mathrm{d}^2\bar{W}/\mathrm{d}y^2 = \mathrm{d}\langle vw \rangle/\mathrm{d}y$, against the DNS profile (the only closure-dependent input of the PINN stage).

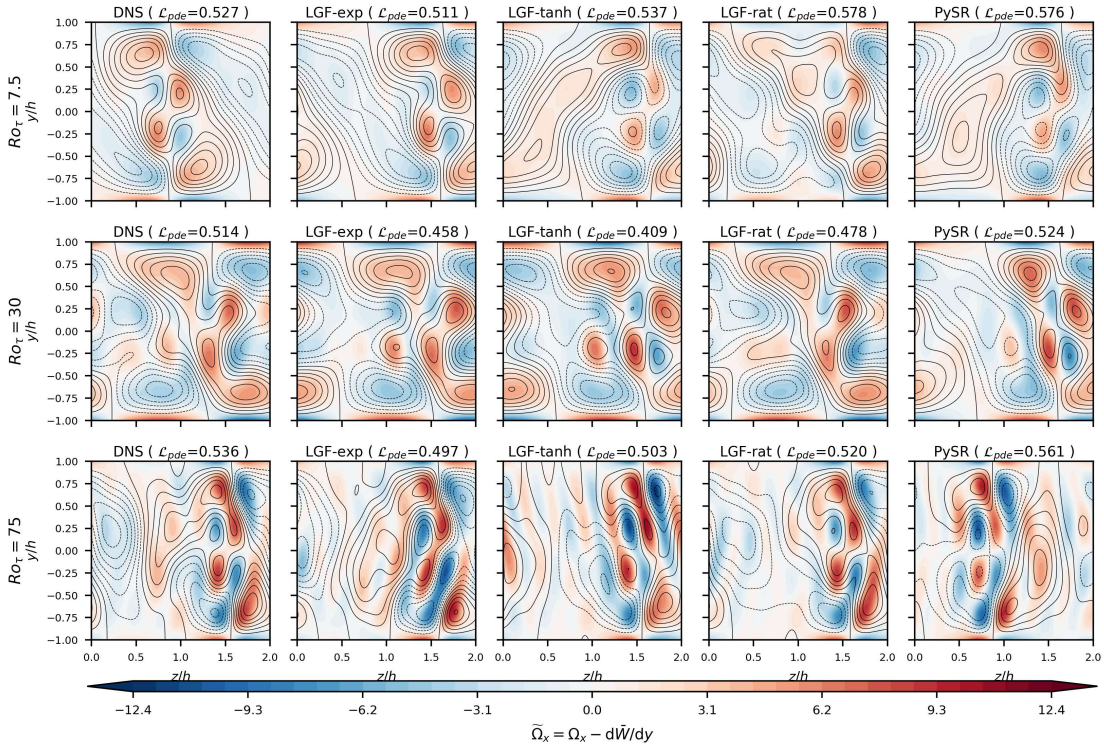


Fig. 6 Secondary flow at three rotation numbers ($Ro_\tau = 7.5$, 30 and 75, one row each) and $\lambda_3 = 2h$: filled contours of the fluctuation vorticity $\tilde{\Omega}_x = \Omega_x - \mathrm{d}\bar{W}/\mathrm{d}y$ with stream-function contour lines (black), for the DNS baseline (left column) and the four closure-driven reconstructions (*exp*, *tanh*, *rat*, PySR), one spanwise wavelength per panel; panel titles give $\mathcal{L}_{\mathrm{pde}}$.

equation and its amplitude from the calibrated energy, robust to O(1) errors in the modelled $\bar{W}$. This robustness makes the two-stage pipeline practical with compact, approximate closures.

The subtraction of $\mathrm{d}\bar{W}/\mathrm{d}y$ in $\tilde{\Omega}_x$ is essential: the total vorticity is dominated by the wall gradient. Two remarks apply. First, the strongest vorticity row flips across runs at $Ro_\tau = 7.5$ (-0.26 for the DNS, *exp* and *rat* against +0.72–0.74 for tanh and PySR), a *solution-branch selection* of the near-symmetric closed equation rather than a model-quality index; settling these branches requires a multi-seed study. Second, the layering fineness drifts with rotation number: the reconstructions bracket the DNS pair count at $Ro_\tau = 30$ (four to eight against six) but overshoot it at $Ro_\tau = 75$ (eight to twelve against six); the finer branches are selected where the modelled $\bar{W}$ misplaces the near-wall vorticity source.

Table 3 Secondary-flow diagnostics at $\lambda_3 = 2h$ for the three simulated rotation numbers. "Pairs" counts the sign changes of $\tilde{\Omega}_x$ per spanwise wavelength ("Row y/h" gives the wall-normal location of the strongest vorticity row); within each rotation block, bold marks the closure value closest to the DNS.

Ro_τ	Source	$\mathcal{L}_{\text{pde}}$	max $\lvert\tilde{v}\rvert$	max $\lvert\tilde{w}\rvert$	max $\lvert\tilde{\Omega}_x\rvert$	Row y/h	Pairs
7.5	DNS	0.527	0.511	0.438	7.92	−0.26	4
	exp	0.511	0.537	**0.455**	**8.28**	−0.25	**4**
	tanh	**0.537**	**0.499**	0.474	8.78	+0.74	2
	rat	0.578	0.493	0.459	8.34	−0.26	6
	PySR	0.576	0.526	0.486	8.93	+0.72	2
30	DNS	0.514	0.620	0.569	10.54	+0.25	6
	exp	0.458	0.582	0.608	10.39	+0.26	4
	tanh	0.409	0.566	0.582	**10.68**	−0.22	**5**
	rat	0.478	0.678	0.589	10.92	+0.25	4
	PySR	**0.524**	**0.585**	**0.574**	9.94	−0.23	8
75	DNS	0.536	0.882	0.458	9.23	+0.69	6
	exp	0.497	0.841	0.496	10.90	−0.24	**8**
	tanh	0.503	0.933	**0.462**	12.40	+0.67	10
	rat	**0.520**	**0.865**	0.507	**9.39**	+0.27	12
	PySR	0.561	0.902	0.477	10.82	+0.26	10

3.4 Discussion

The pipeline divides the labor cleanly: LGF discovers a compact b_{23} closure; the mode-1 balance turns it into the mean spanwise profile $\bar{W}(y)$; the PINN supplies the roll-cell structure and amplitude. The end-to-end error propagation is favorable: even where the modelled $\bar{W}$ departs most from the DNS (0.08–1.06 rms), the reconstructed roll cells stay within 1.9–9.4% in amplitude, because the roll cells are governed by the stress-production and viscous terms of Eq. (2), with the mean-flow advection entering as a correction. Three extensions follow naturally: a sharper near-wall basis would capture more of the steepness of b_{23} that the DNS resolves; a companion closure for $\langle vv \rangle$ would make the amplitude fully autonomous; and a multi-seed study would settle the mild branch selection expected of the near-symmetric equation. These extensions would further strengthen an already practical route: compact learned closures driving a PINN to quantitatively accurate secondary flows at a fraction of the cost of a DNS.

4. Conclusions

A two-stage AI framework has been developed for turbulent secondary flow in the streamwise-rotating channel, combining equation discovery with physics-informed simulation. The LGF engine, an LLM-guided Graph-Feynman search with edge-featured DAG representations, physics priors and an experience-driven outer loop, was set to discover the secondary-flow covariance $b_{23} = \langle vw \rangle / 2k$, the single anisotropy component that drives the problem, from the DNS statistics of five rotation numbers. Across twenty candidates spanning four feature pools (three rotation-variant tensor bases plus their union, each carrying the four symmetry-allowed T_{23} channels), the guided search proved stable rather than merely lucky. The eight LGF runs held the case-normalized error within 0.051–0.137 at 13–17 nodes, whereas the twelve PySR runs spread over 0.042–0.334 with seed-to-seed variation up to a factor of eight and about 60% more nodes. LGF won the g and e pools outright, and its lightly guided go arm beat the fully guided mg arm in every pool, so the LLM's decisive contribution is the physics-informed initialization. The four selected closures (0.042–0.055, five to seven constants each) share one skeleton: a saturating function of the local rotation number, $\sqrt{Ro_{\text{loc}}}$ or $\tanh(cRo_{\text{loc}})$, crossed with a wall function of y_w; one of the two saturation forms recurs in every LGF run, evidence that the search converges to the physics (the balance between Coriolis production and isotropization) rather than to a lucky formula.

Carried through the mean-spanwise-profile balance into a physics-informed spectral network, the four best closures reconstructed the two-layer counter-rotating TG roll cells at $Ro_\tau = 7.5$, 30 and 75 in minutes on a desktop CPU. The secondary-flow amplitudes stay within 1.9–9.4% of the DNS in all twelve runs, and the best matches are structural: the *exp* closure recovers the DNS layering at $Ro_\tau = 7.5$ exactly, the *exp* and *rat* closures keep the DNS row location at $Ro_\tau = 30$, and the tanh closure at $Ro_\tau = 75$ places the strongest row at +0.67 against +0.69 with the amplitude within 6%. The reconstruction is robust to O(1) errors in the modelled mean profile (the smallest and the largest profile errors of the campaign yield comparable amplitude errors). This identifies stress production and viscosity, not mean-flow advection, as the governing balance of the roll cells and makes the two-stage pipeline practical. Natural next steps include a sharper near-wall basis for the steepness of b_{23}, a discovered closure for the amplitude-determining $\langle vv \rangle$, multi-seed confirmation of the branch selection, and the extension of the framework to assimilation of measured profiles and to other rotation-dominated flows.

Conflict of interest

The authors declare that they have no competing interests.

Acknowledgment

This work was supported by the National Natural Science Foundation of China (Grant Nos. U2341242, 12272357).

Data availability

The DNS statistics used in this study are those published by Yang and Wang [3]. The LGF and PINN implementations, together with the prompt templates and run records, are available from the corresponding author on reasonable request.

References

[1] Bradshaw P. Turbulent secondary flows [J]. *Annual Review of Fluid Mechanics*, 1987, 19: 53-74.

[2] Johnston J P, Halleen R M, Lezius D K. Effects of spanwise rotation on the structure of two-dimensional fully developed turbulent channel flow [J]. *Journal of Fluid Mechanics*, 1972, 56: 533-557.

[3] Yang Z, Wang B C. Capturing Taylor–Görtler vortices in a streamwise-rotating channel at very high rotation numbers [J]. *Journal of Fluid Mechanics*, 2018, 838: 658-689.

[4] Speziale C G, Thangam S. Numerical study of secondary flows and roll-cell instabilities in rotating channel flow [J]. *Journal of Fluid Mechanics*, 1983, 130: 377-395.

[5] Masuda S, Fukuda S, Nagata M. Instabilities of plane Poiseuille flow with a streamwise system rotation [J]. *Journal of Fluid Mechanics*, 2008, 603: 189-206.

[6] Wall D P, Nagata M. Nonlinear secondary flow through a rotating channel [J]. *Journal of Fluid Mechanics*, 2006, 564: 25-55.

[7] Kristoffersen R, Andersson H I. Direct simulations of low-Reynolds-number turbulent flow in a rotating channel [J]. *Journal of Fluid Mechanics*, 1993, 256: 163-197.

[8] Wu H, Kasagi N. Effects of arbitrary directional system rotation on turbulent channel flow [J]. *Physics of Fluids*, 2004, 16(4): 979-990.

[9] Oberlack M, Cabot W, Pettersson Reif B A, et al. Group analysis, direct numerical simulation and modelling of a turbulent channel flow with streamwise rotation [J]. *Journal of Fluid Mechanics*, 2006, 562: 383-403.

[10] Yang Y T, Su W D, Wu J Z. Helical-wave decomposition and applications to channel turbulence with streamwise rotation [J]. *Journal of Fluid Mechanics*, 2010, 662: 91-122.

[11] Bongard J, Lipson H. Automated reverse engineering of nonlinear dynamical systems [J]. *Proceedings of the National Academy of Sciences*, 2007, 104(24): 9943-9948.

[12] Schmidt M, Lipson H. Distilling free-form natural laws from experimental data [J]. *Science*, 2009, 324(5923): 81-85.

[13] Brunton S L, Proctor J L, Kutz J N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems [J]. *Proceedings of the National Academy of Sciences*, 2016, 113(15): 3932-3937.

[14] Rudy S H, Brunton S L, Proctor J L, et al. Data-driven discovery of partial differential equations [J]. *Science Advances*, 2017, 3(4): e1602614.

[15] Udrescu S M, Tegmark M. AI Feynman: a physics-inspired method for symbolic regression [J]. *Science Advances*, 2020, 6(16): eaay2631.

[16] Cranmer M. Interpretable machine learning for science with PySR and SymbolicRegression.jl [EB/OL]. 2023 [2026-10-09]. https://arxiv.org/abs/2305.01582.

[17] Shojaee P, Meidani K, Gupta S, et al. LLM-SR: scientific equation discovery via programming with large language models [C]. Proceedings of the 13th International Conference on Learning Representations (ICLR, Oral). 2025.

[18] Xu H, Chen Y, Zhang D. Beyond empirical models: discovering constitutive laws in solids with graph-based equation discovery [J]. *Science Advances*, 2026, 12(37): eaec0989.

[19] Ling J, Kurzawski A, Templeton J. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance [J]. *Journal of Fluid Mechanics*, 2016, 807: 155-166.

[20] Weatheritt J, Sandberg R D. A novel evolutionary algorithm applied to algebraic modifications of the RANS stress–strain relationship [J]. *Journal of Computational Physics*, 2016, 325: 22-37.

[21] Pope S B. A more general effective-viscosity hypothesis [J]. *Journal of Fluid Mechanics*, 1975, 72: 331-340.

[22] Gatski T B, Speziale C G. On explicit algebraic stress models for complex turbulent flows [J]. *Journal of Fluid Mechanics*, 1993, 254: 59-78.

[23] Xiao H, Cinnella P. Quantification of model uncertainty in RANS simulations: a review [J]. *Progress in Aerospace Sciences*, 2019, 108: 1-31.

[24] Duraisamy K, Iaccarino G, Xiao H. Turbulence modeling in the age of data [J]. *Annual Review of Fluid Mechanics*, 2019, 51: 357-377.

[25] 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]. *Journal of Computational Physics*, 2019, 378: 686-707.

[26] Karniadakis G E, Kevrekidis I G, Lu L, et al. Physics-informed machine learning [J]. *Nature Reviews Physics*, 2021, 3: 422-440.

[27] Bai X d, Wang Y, Zhang W. Applying physics informed neural network for flow data assimilation [J]. *Journal of Hydrodynamics*, 2020, 32(6): 1050-1058.

[28] Xu H, Zhang W, Wang Y. Explore missing flow dynamics by physics-informed deep learning: the parameterized governing systems [J]. *Physics of Fluids*, 2021, 33: 095116.

[29] Eivazi H, Tahani M, Schlatter P, et al. Physics-informed neural networks for solving Reynolds-averaged Navier–Stokes equations [J]. *Physics of Fluids*, 2022, 34: 075117.

[30] Wang S, Teng Y, Perdikaris P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks [J]. *SIAM Journal on Scientific Computing*, 2021, 43(5): A3055-A3081.

[31] Wang S, Yu X, Perdikaris P. When and why PINNs fail to train: a neural tangent kernel perspective [J]. *Journal of Computational Physics*, 2022, 449: 110768.

[32] Pope S B. Turbulent flows [M]. Cambridge: Cambridge University Press, 2000.

[33] Launder B E, Reece G J, Rodi W. Progress in the development of a reynolds-stress turbulence closure [J]. *Journal of Fluid Mechanics*, 1975, 68(3): 537-566.

[34] Wallin S, Johansson A V. An explicit algebraic reynolds stress model for incompressible and compressible turbulent flows [J]. *Journal of Fluid Mechanics*, 2000, 403: 89-132.

[35] Romera-Paredes B, Barekatain M, Novikov A, et al. Mathematical discoveries from program search with large language models [J]. *Nature*, 2024, 625(7995): 468-475.