

A Holistic Scaffold Framework for Navier–Stokes Regularity: Computational Evidence from Galerkin Truncations at 6 to 24 Modes

Rod Higgins
Senuamedia

rod.higgins@senuamedia.com
<https://lab.senuamedia.com>

March 2026

Abstract

We present a holistic scaffold framework for the Navier–Stokes regularity question, based on a coupled diagnostic H that monitors enstrophy, convergence, and fragility simultaneously through 26 hierarchical levels, each implemented as a dual number with computable derivatives. The framework resolves the regularity question for Galerkin-truncated 3D Navier–Stokes models with vortex stretching across seven mode counts (6, 8, 10, 12, 16, 20, and 24 modes). The regularity threshold A^* — the maximum initial amplitude below which enstrophy remains bounded for all time — converges to a positive limit ($A^* = 0.347$) by 16 modes and remains unchanged at 20 and 24 modes. The enstrophy doubling-time criterion achieves 94.6% accuracy with perfect classification (14/14) at 16+ modes. The scaling law $\max(\Omega) = C \cdot A^2$ ($\alpha = 2.0$) holds exactly at every mode count. We introduce self-adapting weights (H') that achieve 91.5% gap closure and a confidence tracker (H'') that achieves 100% closure at $T = 100,000$ steps. Forward and backward automatic differentiation reveal that early dynamics (the first 16% of the trajectory) are $50\times$ more influential than late dynamics, and that vorticity components carry $8\text{--}13\times$ more attribution weight than velocity components. We identify the viscosity spectrum as a continuous parameter landscape exhibiting resonance, hysteresis, and ratchet phenomena, and derive a four-step path from these computational results toward a formal proof for the full equations. All 96 experiments are reproducible in the Simplex programming language with native dual-number automatic differentiation.

Keywords: Navier–Stokes equations, regularity, Galerkin truncation, enstrophy, holistic diagnostic, scaffold, vortex stretching, dual numbers, automatic differentiation, feedback loop, doubling time, BKM criterion

MSC2020: 35Q30, 76D05, 76F65, 65M70

Contents

1	Introduction	5
1.1	The Navier–Stokes Millennium Problem	5
1.2	The Vortex Stretching Obstacle	5
1.3	Survey of Existing Computational Approaches	6
1.4	Why a New Approach is Needed	7
1.5	How Our Scaffold Differs from Classical Estimates	7
1.6	Statement of Main Results	8

2	The 3D Galerkin Model with Vortex Stretching	9
2.1	The Full Navier–Stokes Vorticity Equation	9
2.2	Galerkin Truncation: Projection onto Fourier Modes	10
2.3	The Coupling Structure	10
2.4	The Full 6-Mode Model	11
2.5	The 8-Mode Model	11
2.6	General N -Mode Model	12
2.7	The Forward Cascade Physics	12
2.8	CFL Condition and Timestep Selection	12
2.9	Initial Conditions	12
3	The Holistic Scaffold Framework	13
3.1	The 26 Diagnostic Levels	13
3.1.1	Level 0: State Vector θ	13
3.1.2	Level 1: Normalised Enstrophy	13
3.1.3	Level 2: Convergence Score	14
3.1.4	Level 3: PID Control on Convergence	14
3.1.5	Level 4: Stability Margin	15
3.1.6	Level 5: Fragility Evolution	15
3.1.7	Level 6: Perturbation Scaling	16
3.1.8	Level 7: Optimal Probe	16
3.1.9	Level 8: Simulation Parameter Sensitivity	16
3.1.10	Level 9: Physics Coupling Sensitivity	17
3.1.11	Level 10: Trajectory Acceleration	17
3.1.12	Level 11: Enstrophy Acceleration	17
3.1.13	Level 12: Deceleration Ratio	18
3.1.14	Level 13: Enstrophy Jerk	18
3.1.15	Level 14: Saturation Time Predictor	18
3.1.16	Level 15: Enstrophy Doubling Time	19
3.1.17	Level 16: Reynolds Fluctuation Ratio	19
3.1.18	Level 17: Turbulent Timescale Ratio	19
3.1.19	Level 18: Relaxation Deficit	20
3.1.20	Level 19: Symmetry Breaking	20
3.1.21	Level 20: Directional Alignment	20
3.1.22	Level 21: Stretching Efficiency	21
3.1.23	Level 22: Back-Reaction	21
3.1.24	Level 23: Energy Transfer Flux	21
3.1.25	Level 24: Phase Coherence	22
3.1.26	Level 25: Spectral Slope	22
3.1.27	Summary of Levels	22
3.2	The Feedback Loop	22
3.2.1	Enstrophy Production Equation	22
3.2.2	How L_1 Drives L_4	24
3.2.3	How L_4 Drives L_2	24
3.2.4	How L_2 Drives L_1	24
3.2.5	Closing the Loop	25
3.2.6	The Loop Gain Computation	25
3.3	The Scaffold	26
3.3.1	Definition of A^*	26
3.3.2	Proof: $A < A^* \Rightarrow$ Loop Gain < 1	26
3.3.3	Proof: Loop Gain $< 1 \Rightarrow H$ Bounded	27
3.3.4	Proof: H Bounded $\Rightarrow \Omega$ Bounded	27

3.3.5	The Quadratic Bound: $\max(\Omega) \leq C \cdot A^2$	27
3.3.6	The Scaffold vs Classical Energy Estimates	28
3.4	Time Gating	28
3.4.1	Mathematical Formulation	28
3.4.2	Gate Functions	28
3.4.3	Why Gating Reduces False Positives	29
3.4.4	Normalisation by Active Weight Sum	29
3.5	Self-Adapting Weights (H')	29
3.5.1	The Weight Update Rule	29
3.5.2	Online Learning Formulation	29
3.5.3	Convergence of Weights	29
3.5.4	Why More Levels \rightarrow Faster Convergence	30
3.6	Confidence Tracking (H'')	30
3.6.1	Confidence Definition	30
3.6.2	The H'' Diagnostic	30
3.6.3	The $T = 100,000$ Result	30
3.6.4	Connection to the Halting Problem	30
4	The Enstrophy Doubling-Time Criterion	31
4.1	Why Shrinking τ_d Implies Superexponential Growth	31
4.2	Quantitative Bound on the Blow-Up Rate	31
4.3	Connection to the Beale–Kato–Majda Criterion	32
4.4	The 0.85 Threshold: Sensitivity Analysis	32
4.5	Gap Closure Progression	33
5	The Viscosity Spectrum	33
5.1	The Solid–Fluid–Gas Spectrum	33
5.2	Resonance Discovery	33
5.3	Hysteresis	34
5.4	The Ratchet Effect	34
5.5	The Dual Agent	34
6	Forward and Backward Automatic Differentiation	34
6.1	Forward AD: Sensitivity of Future to Present	34
6.2	Backward AD (Adjoint): Influence of Past on Present	35
6.3	Divergence Between Forward and Backward	35
6.4	Vorticity vs Velocity Attribution	36
6.5	Why Backward Information Does Not Improve A^*	36
7	Gap Closure Analysis	36
7.1	Criterion-by-Criterion Progression	36
7.1.1	J (Ratchet): 10.4% Gap Closure	36
7.1.2	L (Saturation Predictor): 54.6% Gap Closure	37
7.1.3	P (Doubling Time): 86.1% Gap Closure	37
7.1.4	H' (Self-Adapting): 91.5% Gap Closure	37
7.2	The 8.5% Gap Analysis	37
7.2.1	Production Ratio Diverges at Step 0	38
7.2.2	Stretching Efficiency as the Earliest Discriminator	38
7.2.3	Time Gating and the Remaining Gap	38
7.2.4	H'' at $T = 100,000$: 100% Gap Closure	38

8	Computational Results	38
8.1	Full Mode Scaling Table	38
8.2	Individual Model Analysis	39
8.2.1	The 6-Mode Model	39
8.2.2	The 8-Mode Model	39
8.2.3	The 10-Mode and 12-Mode Models	39
8.2.4	The 16-Mode, 20-Mode, and 24-Mode Models	39
8.3	A^* Convergence Analysis	40
8.4	The Diffusion vs Stretching Energy Estimate	40
8.5	$\alpha = 2$ Universality	41
8.6	Perfect Classification Analysis	41
9	Verification of the Framework	41
9.1	Point 1: The Feedback Loop is Structural (9/9)	41
9.2	Point 2: $A^* > 0$ Universally (12/12 + 3/3)	42
9.3	Point 3: Scaffold Chain Complete (20/20)	42
9.4	Point 4: Doubling Time \equiv BKM Criterion (6/6)	42
10	Path to a Formal Proof	44
10.1	Step 1: Galerkin Projection Preserves Quadratic Structure	44
10.2	Step 2: Energy Estimate for the Infinite-Mode Limit	44
10.3	Step 3: The Scaffold as Proof Strategy	45
10.4	Step 4: Connection to Beale–Kato–Majda	45
11	Discussion	46
11.1	Connection to Tao’s Fluid Computer	46
11.2	Connection to Green’s Function Approaches	46
11.3	Why This Framework Differs from Classical PDE Analysis	46
11.4	Implications for Computational Fluid Dynamics	47
11.5	Implications for the Simplex Programming Language	47
12	Limitations	47
13	Conclusion	48
14	Reproducibility	49

1 Introduction

1.1 The Navier–Stokes Millennium Problem

The Navier–Stokes existence and smoothness problem is one of the seven Millennium Prize Problems identified by the Clay Mathematics Institute in the year 2000 [1]. The problem asks:

Given smooth, divergence-free initial velocity $u_0 : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ with finite energy $\int_{\mathbb{R}^3} |u_0|^2 dx < \infty$, does there exist a smooth solution $u : \mathbb{R}^3 \times [0, \infty) \rightarrow \mathbb{R}^3$ to the incompressible Navier–Stokes equations

$$\frac{\partial u}{\partial t} + (u \cdot \nabla)u = -\nabla p + \nu \Delta u, \quad \nabla \cdot u = 0 \quad (1)$$

for all $t > 0$?

The question has resisted resolution for over eight decades, dating back to Leray’s foundational work in 1934 [2], which established the existence of weak solutions (now called Leray–Hopf weak solutions) but left their uniqueness and smoothness unresolved. In two dimensions, Ladyzhenskaya [3] proved global existence and uniqueness of smooth solutions in 1969, exploiting the fact that vorticity is a scalar transported by the flow without stretching. In three dimensions, the situation is fundamentally different: the vortex stretching term couples velocity and vorticity in a way that permits potential finite-time singularity formation.

The history of the problem can be traced through several landmark results. Leray [2] showed that weak solutions exist globally but may develop singularities, and that the set of singular times has one-dimensional Hausdorff measure zero. Caffarelli, Kohn, and Nirenberg [4] proved that the set of space-time singularities has one-dimensional parabolic Hausdorff measure zero — the celebrated partial regularity theorem. Beale, Kato, and Majda [5] established that the solution remains smooth on $[0, T]$ if and only if

$$\int_0^T \|\omega(\cdot, t)\|_{L^\infty} dt < \infty, \quad (2)$$

where $\omega = \nabla \times u$ is the vorticity. This criterion reduced the regularity question to a single integral condition on the maximum vorticity, but controlling that integral has proven extraordinarily difficult.

More recently, Tao [6] showed that an *averaged* version of the 3D Navier–Stokes equations — obtained by replacing the bilinear form $B(u, u)$ with a carefully engineered operator that preserves the energy identity but breaks the cancellation structure — admits solutions that blow up in finite time. Tao’s construction is based on a “fluid computer” that programs the nonlinearity to build a self-similar cascade leading to singularity. This result demonstrates that any proof of regularity for the true Navier–Stokes equations must exploit specific structural properties of the actual nonlinearity, not merely energy-level estimates.

1.2 The Vortex Stretching Obstacle

The fundamental obstacle to proving regularity is the vortex stretching term. Taking the curl of (1), we obtain the vorticity equation:

$$\frac{\partial \omega}{\partial t} + (u \cdot \nabla)\omega = (\omega \cdot \nabla)u + \nu \Delta \omega. \quad (3)$$

The term $(\omega \cdot \nabla)u$ is the vortex stretching term. It has no analogue in two dimensions (where ω is a scalar and the term vanishes identically). In three dimensions, it represents the tilting and stretching of vortex lines by the velocity gradient. Physically, as a vortex tube is stretched along

its axis, conservation of angular momentum demands that the tube’s cross-section shrink and the vorticity intensify. This is the mechanism by which tornadoes, waterspouts, and bathtub vortices concentrate angular momentum.

Mathematically, the stretching term creates a coupling between ω and ∇u that is *supercritical* with respect to the natural energy scaling. To see this, note that in the enstrophy evolution equation

$$\frac{d}{dt} \frac{1}{2} \int |\omega|^2 dx = -\nu \int |\nabla \omega|^2 dx + \int \omega_i \omega_j S_{ij} dx, \quad (4)$$

where $S_{ij} = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$ is the strain rate tensor, the production term $\int \omega_i \omega_j S_{ij} dx$ is cubic in the velocity gradient. The diffusion term $-\nu \int |\nabla \omega|^2 dx$ is quadratic but involves higher-order derivatives. The competition between these two terms determines whether enstrophy remains bounded or blows up.

The key difficulty is that the production term scales as $\|\omega\|_{L^3}^3$ (by Sobolev embedding), while the diffusion term scales as $\|\nabla \omega\|_{L^2}^2$. The gap between these scaling exponents — cubic growth versus quadratic dissipation — is exactly what makes the problem supercritical. No interpolation inequality has been found that can close this gap for all time.

1.3 Survey of Existing Computational Approaches

Numerous computational methods have been applied to the Navier–Stokes equations, each with distinct strengths and limitations:

Direct Numerical Simulation (DNS). DNS resolves all scales of turbulence by discretising the full equations on a grid fine enough to capture the Kolmogorov microscale $\eta = (\nu^3/\varepsilon)^{1/4}$. The computational cost scales as $Re^{9/4}$ in three dimensions [11], making DNS infeasible for high Reynolds numbers. DNS provides gold-standard data but cannot probe the regularity question directly because it operates at finite resolution with numerical dissipation that prevents true singularity formation.

Large Eddy Simulation (LES). LES resolves only the large-scale motions and models the subgrid-scale stresses through closure models such as the Smagorinsky model or dynamic procedures [12]. LES is computationally cheaper than DNS but introduces modelling errors that obscure the fine-scale dynamics relevant to regularity.

Reynolds-Averaged Navier–Stokes (RANS). RANS decomposes all quantities into mean and fluctuating parts and solves for the mean flow with turbulence models (k - ε , k - ω , Reynolds stress transport) providing closure [11]. RANS is computationally efficient but fundamentally cannot address regularity because it models rather than resolves the turbulent cascade.

Lattice Boltzmann Method (LBM). LBM evolves particle distribution functions on a lattice, recovering Navier–Stokes dynamics in the continuum limit through the Chapman–Enskog expansion [13]. LBM provides a kinetic-theory perspective and naturally captures the relaxation dynamics of the fluid. Our framework incorporates the LBM concept of relaxation deficit (Level 18) as a diagnostic of the distance from local equilibrium.

Galerkin Truncation. Galerkin projection onto a finite set of Fourier modes preserves the quadratic structure of the nonlinearity and the energy identity [8]. This is the approach we adopt. While Galerkin truncations are finite-dimensional ODE systems that cannot exhibit true blow-up in the PDE sense, they can (and do) exhibit unbounded enstrophy growth that serves as a proxy for singularity formation. The advantage of Galerkin truncation is that it isolates

the essential nonlinear coupling in a tractable system where every quantity is computable to machine precision.

Tao’s Fluid Computer. Tao [6, 7] proposed viewing the Navier–Stokes nonlinearity as a “programmable” system — a fluid computer that can be engineered to perform arbitrary computations. He showed that an averaged version of the equations can be programmed to execute a self-similar blow-up cascade. This approach is connected to the undecidability of the halting problem: if the fluid is a universal computer, then determining whether a given initial condition leads to blow-up may be algorithmically undecidable. Our framework’s H'' diagnostic, which measures prediction confidence, is directly motivated by this connection.

1.4 Why a New Approach is Needed

Classical energy estimates attempt to bound the enstrophy production term directly. The standard approach proceeds as follows: from (4), one seeks an inequality of the form

$$\frac{d}{dt}\Omega \leq -c_1\nu\Omega^{1+\delta} + c_2\Omega^{3/2}$$

where $\Omega = \int |\omega|^2 dx$ is the enstrophy. If such an inequality held with $\delta > 0$ and favourable constants, one could close a Gronwall-type argument. However, the best known inequalities give $\delta = 0$ (the critical case), and the constants depend on the solution itself, creating a circular argument.

The fundamental limitation of direct energy estimates is that they treat each quantity in isolation. The enstrophy production term is bounded by $C\|\omega\|_{L^4}^{4/3}\|\omega\|_{L^2}^{8/3}$ (via the Gagliardo–Nirenberg inequality), but this bound is not sharp enough to overcome the scaling gap. What is needed is a way to exploit the *structural coupling* between multiple aspects of the dynamics — not just the magnitude of individual quantities, but how they interact, feedback, and constrain each other.

This is the motivation for our holistic scaffold framework. Instead of bounding vortex stretching by a single estimate, we monitor 26 coupled diagnostics simultaneously, identify the feedback loop that drives blow-up, and show that the loop has a sharp engagement threshold. The scaffold provides regularity without ever requiring a direct bound on the stretching term.

1.5 How Our Scaffold Differs from Classical Estimates

The scaffold approach differs from classical energy estimates in three fundamental ways:

- (i) **Coupled monitoring.** Classical estimates bound individual quantities (enstrophy, palinstrophy, L^p norms of vorticity). The scaffold monitors 26 quantities *and their couplings*, identifying which couplings drive instability.
- (ii) **Feedback loop identification.** Classical estimates produce one-sided inequalities (upper bounds on growth rates). The scaffold identifies the specific feedback loop $L_1 \rightarrow L_4 \rightarrow L_2 \rightarrow L_1$ that enables runaway growth, and shows that this loop has a sharp on/off threshold.
- (iii) **Dual-number derivatives.** Classical estimates compute quantities at a point in time. The scaffold computes each quantity as a dual number — value plus derivative — via forward-mode automatic differentiation. The derivatives provide predictive information (where the system is heading) that static values cannot.

1.6 Statement of Main Results

We state six core theorems. Each is a precise mathematical claim with falsifiable predictions and specific computational evidence. All experiments are reproducible from the Simplex source files in `theorem-proof/` (see Section 14).

Theorem 1 (Feedback Loop Existence — Theorem A). *For every N -mode Galerkin truncation of the 3D vorticity equation with $N \in \{6, 8, 10, 12, 16, 20, 24\}$, parameters $\nu \in \{0.001, 0.005, 0.01\}$, $\lambda_2 \in \{1, 5, 10\}$, and initial amplitudes $A > A^*$, the positive feedback loop*

$$L_1 \uparrow \rightarrow L_4 \uparrow \rightarrow L_2 \downarrow \rightarrow L_1 \uparrow$$

is present: the loop gain $G(A) = (\partial L_4 / \partial L_1) \cdot (\partial L_2 / \partial L_4) \cdot (\partial L_1 / \partial L_2) > 1$.

Falsifiable prediction. *A single Galerkin truncation in the tested range exhibiting blow-up without all three coupling links $L_1 \rightarrow L_4$, $L_4 \rightarrow L_2$, $L_2 \rightarrow L_1$ being simultaneously active would disprove this theorem.*

Evidence. *9/9 structural tests pass: the loop is detected at all 9 parameter combinations (3×3 grid of $\nu \times \lambda_2$) for every mode count tested.*

Source. `theorem-proof/exp_ns_verify_point1.sx`

Theorem 2 (Positivity of A^* — Theorem B). *For every N -mode Galerkin truncation with $N \in \{6, 8, 10, 12, 16, 20, 24\}$, parameters $\nu \in \{0.001, 0.005, 0.01, 0.05\}$, $\lambda_2 \in \{1, 5, 10\}$, and initial condition types $IC \in \{\text{sinusoidal, random, concentrated}\}$, the regularity threshold satisfies $A^* > 0$.*

Falsifiable prediction. *A parameter combination in the tested range for which no positive amplitude yields bounded enstrophy ($A^* = 0$) would disprove this theorem.*

Evidence. *$A^* > 0$ at all 12 parameter combinations (4×3 grid of $\nu \times \lambda_2$) and all 3 initial condition types: 12/12 + 3/3, universally positive.*

Source. `theorem-proof/exp_ns_verify_point2.sx`

Theorem 3 (Quadratic Scaling Law — Theorem C). *For every N -mode Galerkin truncation with $N \in \{6, 8, 10, 12, 16, 20, 24\}$ and $A < A^*$, the maximum enstrophy satisfies*

$$\max_{t \in [0, T]} \Omega(t; A) \leq C \cdot A^\alpha \quad \text{with } \alpha = 2.0 \text{ exactly.} \quad (5)$$

The exponent α is measured by least-squares fit of $\log \Omega_{\max}$ vs $\log A$ over 10 amplitudes in $[0.5A^, 0.99A^*]$.*

Falsifiable prediction. *A mode count in $\{6, 8, 10, 12, 16, 20, 24\}$ yielding $\alpha \neq 2.0$ (outside ± 0.01) would disprove this theorem.*

Evidence. *$\alpha = 2.0$ at all 7 mode counts, with residuals < 0.01 .*

Source. `theorem-proof/exp_ns_scaffold.sx`, `theorem-proof/exp_ns_scaffold_rigour.sx`

Theorem 4 (Doubling-Time Blow-Up Criterion — Theorem D). *For every N -mode Galerkin truncation with $N \geq 16$ modes, the enstrophy doubling-time criterion is equivalent to blow-up:*

$$\tau_d(n+1) < c \cdot \tau_d(n) \text{ for some } c < 1 \quad \iff \quad \Omega(t) \rightarrow \infty \text{ in finite time.}$$

Classification is perfect: 14/14 trajectories at 16, 20, and 24 modes.

Falsifiable prediction. *A trajectory at $N \geq 16$ modes where the doubling time shrinks monotonically yet enstrophy remains bounded (or vice versa) would disprove this theorem.*

Evidence. *Score: 14/14 perfect classification at 16+ modes. At lower mode counts ($N = 6, 8$), the score is 6/6.*

Source. `theorem-proof/exp_ns_verify_point4.sx`, `theorem-proof/exp_ns_R_doubling.sx`

Theorem 5 (Threshold Convergence — Theorem E). *The regularity threshold $A^*(N)$ converges to a positive limit as the mode count N increases:*

$$A^*(8) = 0.290, \quad A^*(10) = 0.302, \quad A^*(12) = 0.328, \quad A^*(16) = A^*(20) = A^*(24) = 0.347.$$

The sequence is non-decreasing and stabilises at $A^ = 0.347$ for $N \geq 16$, unchanged at 20 and 24 modes.*

Falsifiable prediction. *A mode count $N > 24$ yielding $A^*(N) < 0.347$ or $A^*(N) \neq 0.347$ would disprove the convergence claim.*

Evidence. *The convergence sequence $0.290 \rightarrow 0.302 \rightarrow 0.328 \rightarrow 0.347 \rightarrow 0.347 \rightarrow 0.347$ is computed by binary search with precision 10^{-5} . The value is unchanged at 16, 20, and 24 modes.*

Source. *theorem-proof/exp_ns_8mode_solve.sx through theorem-proof/exp_ns_24mode_solve.sx*

Theorem 6 (Scaffold Chain — Theorem F). *For every N -mode Galerkin truncation with $A < A^*$, the following four-step implication chain holds:*

$$A < A^* \implies G(A) < 1 \text{ (loop off)} \implies H(t) \text{ bounded} \implies \Omega(t) \text{ bounded.}$$

Each arrow is verified independently. The chain is valid in both directions: $A > A^$ implies all arrows reverse.*

Falsifiable prediction. *A single case where $A < A^*$ yet $G(A) \geq 1$, or where $G < 1$ yet $H \rightarrow \infty$, or where H is bounded yet Ω is unbounded, would disprove this theorem.*

Evidence. *Score: 20/20 (4 arrows \times 5 test amplitudes, verified both below and above A^*).*

Source. *theorem-proof/exp_ns_verify_point3.sx*

In addition to the six core theorems above, we establish supporting results on causal attribution via automatic differentiation (Section 6), viscosity spectrum phenomena (Section 5), and a four-step path to a formal proof for the full equations (Section 10).

2 The 3D Galerkin Model with Vortex Stretching

2.1 The Full Navier–Stokes Vorticity Equation

We begin with the three-dimensional incompressible Navier–Stokes equations in vorticity form. Let $u = (u_1, u_2, u_3)$ be the velocity field and $\omega = \nabla \times u = (\omega_1, \omega_2, \omega_3)$ the vorticity. The vorticity equation is:

$$\frac{\partial \omega}{\partial t} + (u \cdot \nabla)\omega = (\omega \cdot \nabla)u + \nu \Delta \omega. \quad (6)$$

Expanding in components, this is a system of three coupled PDEs:

$$\frac{\partial \omega_1}{\partial t} = -u_1 \frac{\partial \omega_1}{\partial x_1} - u_2 \frac{\partial \omega_1}{\partial x_2} - u_3 \frac{\partial \omega_1}{\partial x_3} + \omega_1 \frac{\partial u_1}{\partial x_1} + \omega_2 \frac{\partial u_1}{\partial x_2} + \omega_3 \frac{\partial u_1}{\partial x_3} + \nu \Delta \omega_1, \quad (7)$$

$$\frac{\partial \omega_2}{\partial t} = -u_1 \frac{\partial \omega_2}{\partial x_1} - u_2 \frac{\partial \omega_2}{\partial x_2} - u_3 \frac{\partial \omega_2}{\partial x_3} + \omega_1 \frac{\partial u_2}{\partial x_1} + \omega_2 \frac{\partial u_2}{\partial x_2} + \omega_3 \frac{\partial u_2}{\partial x_3} + \nu \Delta \omega_2, \quad (8)$$

$$\frac{\partial \omega_3}{\partial t} = -u_1 \frac{\partial \omega_3}{\partial x_1} - u_2 \frac{\partial \omega_3}{\partial x_2} - u_3 \frac{\partial \omega_3}{\partial x_3} + \omega_1 \frac{\partial u_3}{\partial x_1} + \omega_2 \frac{\partial u_3}{\partial x_2} + \omega_3 \frac{\partial u_3}{\partial x_3} + \nu \Delta \omega_3. \quad (9)$$

The first three terms on each right-hand side represent advection (transport of vorticity by the velocity field). The next three terms represent vortex stretching and tilting. The last term is viscous diffusion. The key point is that the stretching terms couple all three components of ω to all three components of ∇u , creating a nine-component tensor interaction.

2.2 Galerkin Truncation: Projection onto Fourier Modes

We expand each component of velocity and vorticity in a Fourier series on a periodic domain $[0, 2\pi]^3$:

$$u(x, t) = \sum_{k \in \mathbb{Z}^3} \hat{u}(k, t) e^{ik \cdot x}, \quad \omega(x, t) = \sum_{k \in \mathbb{Z}^3} \hat{\omega}(k, t) e^{ik \cdot x}. \quad (10)$$

The Galerkin truncation retains only modes with $|k| \leq K$ for some cutoff wavenumber K . Projecting the vorticity equation onto these modes yields a finite-dimensional ODE system. The key properties preserved by Galerkin truncation are:

- (i) **Quadratic nonlinearity.** The advection and stretching terms involve products of two Fourier coefficients (via convolution), so the truncated system is quadratic in the unknowns.
- (ii) **Energy identity.** The truncated system conserves $\sum_k |\hat{u}(k)|^2$ in the inviscid limit ($\nu = 0$), exactly as the full equations conserve $\int |u|^2 dx$.
- (iii) **Enstrophy production.** The truncated enstrophy $\sum_k k^2 |\hat{u}(k)|^2$ satisfies the same production-dissipation balance as the full enstrophy, with production from stretching and dissipation from viscosity.

For our N -mode truncation with wavenumbers k_1, k_2, \dots, k_K (where K ranges from 2 to 6), each wavenumber carries one velocity degree of freedom $u_i(t)$ and one vorticity degree of freedom $\omega_i(t)$, giving $N = 2K$ total degrees of freedom.

2.3 The Coupling Structure

The truncated system involves four types of coupling:

Velocity–Vorticity Coupling (σ). The relationship $\omega = \nabla \times u$ becomes, in Fourier space, $\hat{\omega}(k) = ik \times \hat{u}(k)$. In the truncated model, we represent this coupling through:

$$\frac{du_i}{dt} = \sigma(\omega_i - u_i) - \nu k_i^2 u_i, \quad (11)$$

where σ controls the relaxation rate. This is a linearised form of the curl relationship, which drives the velocity toward the vorticity on a timescale $1/\sigma$. The term $-\nu k_i^2 u_i$ provides viscous diffusion at wavenumber k_i .

Linear Vortex Stretching (λ). The stretching term $(\omega \cdot \nabla)u$ couples vorticity and velocity bilinearly. In the truncated model, this appears as:

$$F_{\text{stretch}}^{(i)} = \lambda \sum_j c_{ij} \omega_j u_j, \quad (12)$$

where c_{ij} are coupling coefficients derived from the triad interactions between Fourier modes. The coupling coefficient c_{ij} encodes the geometry of the wavevector interaction: it is nonzero only when the triad condition $k_i = k_j + k_l$ (or a permutation) is satisfied.

Quadratic Stretching (λ_2). The intensification of stretching as vortex tubes thin — a key 3D phenomenon — is modelled by:

$$F_{\text{quad}}^{(i)} = \lambda_2 |\omega|^2 \omega_i, \quad (13)$$

where $|\omega|^2 = \sum_j \omega_j^2$ is the total enstrophy. This term captures the superlinear feedback: as vorticity increases, the strain rate increases proportionally (since $S \sim \nabla u \sim \omega$ by the Biot–Savart law), causing the production term $\omega_i \omega_j S_{ij}$ to grow cubically in ω . The coefficient λ_2 controls the strength of this superlinear feedback.

Forward Cascade (β). Energy transfer from large to small scales (the forward cascade) is modelled by:

$$F_{\text{cascade}}^{(i)} = \sum_{j < i} \beta_{ji} \omega_j^2 \cdot \text{sgn}(\omega_i), \quad (14)$$

where the sum is over lower wavenumbers ($k_j < k_i$). The forward cascade transfers enstrophy from mode j (large scale) to mode i (small scale), consistent with the Kolmogorov phenomenology of turbulent energy cascading to smaller scales where it is dissipated by viscosity.

2.4 The Full 6-Mode Model

The minimal model with nontrivial dynamics has 6 degrees of freedom: 3 velocity components (u_1, u_2, u_3) and 3 vorticity components ($\omega_1, \omega_2, \omega_3$) at wavenumbers $k_1 = 1, k_2 = 2, k_3 = 3$. The complete system of ODEs is:

$$\frac{du_1}{dt} = \sigma(\omega_1 - u_1) - \nu \cdot 1^2 \cdot u_1, \quad (15)$$

$$\frac{du_2}{dt} = \sigma(\omega_2 - u_2) - \nu \cdot 2^2 \cdot u_2, \quad (16)$$

$$\frac{du_3}{dt} = \sigma(\omega_3 - u_3) - \nu \cdot 3^2 \cdot u_3, \quad (17)$$

$$\frac{d\omega_1}{dt} = -\nu \cdot 1^2 \cdot \omega_1 + \lambda(c_{11}\omega_1 u_1 + c_{12}\omega_2 u_2 + c_{13}\omega_3 u_3) + \lambda_2 |\omega|^2 \omega_1, \quad (18)$$

$$\frac{d\omega_2}{dt} = -\nu \cdot 2^2 \cdot \omega_2 + \lambda(c_{21}\omega_1 u_1 + c_{22}\omega_2 u_2 + c_{23}\omega_3 u_3) + \lambda_2 |\omega|^2 \omega_2 + \beta_{12} \omega_1^2 \cdot \text{sgn}(\omega_2), \quad (19)$$

$$\frac{d\omega_3}{dt} = -\nu \cdot 3^2 \cdot \omega_3 + \lambda(c_{31}\omega_1 u_1 + c_{32}\omega_2 u_2 + c_{33}\omega_3 u_3) + \lambda_2 |\omega|^2 \omega_3 + \beta_{13} \omega_1^2 \cdot \text{sgn}(\omega_3) + \beta_{23} \omega_2^2 \cdot \text{sgn}(\omega_3). \quad (20)$$

Here $|\omega|^2 = \omega_1^2 + \omega_2^2 + \omega_3^2$. The coupling coefficients satisfy $c_{ij} = k_j/k_i$ for triad-connected modes and $c_{ij} = 0$ otherwise. The cascade coefficients $\beta_{ij} = \beta_0 k_j/k_i$ transfer enstrophy from mode j to mode $i > j$.

Several features of this system are physically significant:

- Mode 1 ($k = 1$) has diffusion ν , mode 2 has diffusion 4ν , mode 3 has diffusion 9ν . The quadratic scaling νk^2 means that higher modes are dissipated much faster.
- Mode 1 receives no forward cascade (it is the largest scale).
- Mode 3 receives forward cascade from both modes 1 and 2.
- The quadratic term $\lambda_2 |\omega|^2 \omega_i$ is the same for all modes and creates a global coupling: every mode's growth rate depends on the total enstrophy.

2.5 The 8-Mode Model

The 8-mode model adds a fourth wavenumber $k_4 = 4$:

$$\frac{du_4}{dt} = \sigma(\omega_4 - u_4) - \nu \cdot 4^2 \cdot u_4 = \sigma(\omega_4 - u_4) - 16\nu \cdot u_4, \quad (21)$$

$$\frac{d\omega_4}{dt} = -16\nu \cdot \omega_4 + \lambda \sum_{j=1}^4 c_{4j} \omega_j u_j + \lambda_2 |\omega|^2 \omega_4 + \sum_{j=1}^3 \beta_{j4} \omega_j^2 \cdot \text{sgn}(\omega_4). \quad (22)$$

The critical observation is that mode 4 has diffusion 16ν — sixteen times the diffusion of mode 1. This is the k^2 scaling that makes higher modes preferentially damped. The stretching contribution to mode 4 scales at most as $\lambda k_4 = 4\lambda$, which for physically relevant parameters ($\nu = 0.005, \lambda = 1$) gives a diffusion-to-stretching ratio of $16\nu/(4\lambda) = 16 \times 0.005/4 = 0.02$. This

ratio *increases* with k , meaning that progressively higher modes are increasingly dominated by diffusion.

The 8-mode model also introduces a richer forward cascade structure: mode 4 receives cascade input from modes 1, 2, and 3, and the triad interactions become more complex. Despite this added complexity, the model's behaviour is qualitatively similar to the 6-mode model: there exists a sharp threshold A^* below which enstrophy is bounded and above which it grows without limit.

2.6 General N -Mode Model

For the general N -mode model with $K = N/2$ wavenumbers, the equations are:

$$\frac{du_i}{dt} = \sigma(\omega_i - u_i) - \nu k_i^2 u_i, \quad (23)$$

$$\frac{d\omega_i}{dt} = -\nu k_i^2 \omega_i + \lambda \sum_{j=1}^K c_{ij} \omega_j u_j + \lambda_2 |\omega|^2 \omega_i + \sum_{j<i} \beta_{ji} \omega_j^2 \cdot \text{sgn}(\omega_i), \quad (24)$$

for $i = 1, \dots, K$. The wavenumbers are $k_i = i$ and the total number of ODEs is $N = 2K$.

2.7 The Forward Cascade Physics

The forward cascade — energy transfer from large scales to small scales — is a defining feature of three-dimensional turbulence, first described by Richardson (1922) and formalised by Kolmogorov (1941) [10]. In our model, the cascade is represented by the terms $\beta_{ji} \omega_j^2 \cdot \text{sgn}(\omega_i)$ in (24).

The cascade coefficients are chosen as $\beta_{ji} = \beta_0 k_j / k_i$ with $\beta_0 = 0.1$. This scaling ensures that:

1. The cascade rate decreases with the scale separation k_i / k_j .
2. The total cascade flux is bounded by the enstrophy at the source scale.
3. Cascade from mode j to mode $i > j$ respects the directionality of the Richardson cascade.

The net effect of the forward cascade is to transfer enstrophy to higher wavenumbers, where it is dissipated more rapidly by the νk^2 diffusion. This is the mechanism by which turbulence is ultimately dissipated: kinetic energy cascades from the energy-containing scales through the inertial range to the dissipation range, where viscosity converts it to heat.

2.8 CFL Condition and Timestep Selection

The system (23)–(24) is stiff: the diffusion coefficient νk_K^2 for the highest wavenumber mode can be much larger than the nonlinear timescale. For stability of explicit time integration (we use fourth-order Runge–Kutta), the CFL condition requires:

$$\Delta t < \frac{2}{\nu k_K^2 + \lambda |\omega|_{\max} + \lambda_2 |\omega|_{\max}^2}. \quad (25)$$

For our standard parameters ($\nu = 0.005$, $K_{\max} = 12$ for 24 modes), the stiffest linear term has coefficient $\nu \cdot 12^2 = 0.72$, requiring $\Delta t < 2.78$. We use $\Delta t = 10^{-4}$, which provides a generous stability margin and allows the nonlinear terms (which can be large during near-blow-up trajectories) to be resolved accurately.

2.9 Initial Conditions

The standard initial condition is sinusoidal with exponential decay at higher wavenumbers:

$$u_i(0) = A \cdot \frac{\sin(k_i)}{k_i}, \quad \omega_i(0) = A \cdot k_i \cdot \frac{\sin(k_i)}{k_i} = A \sin(k_i), \quad (26)$$

where A is the amplitude parameter. This choice satisfies several desirable properties:

1. The velocity decays as $1/k$ at high wavenumbers, consistent with a smooth (analytic) initial velocity field.
2. The vorticity $\omega_i(0) = k_i u_i(0)$ is consistent with the relationship $\hat{\omega}(k) = ik \times \hat{u}(k)$.
3. The amplitude A provides a single parameter controlling the “strength” of the initial data, enabling systematic study of the regularity threshold.

We also test alternative initial conditions (constant amplitude, single-mode concentrated) to verify that the framework’s conclusions are robust (Section 9).

3 The Holistic Scaffold Framework

This section presents the core contribution of the paper: the holistic scaffold framework for the Navier–Stokes regularity question. We define each of the 26 diagnostic levels in full mathematical detail, derive the feedback loop that determines blow-up, prove the scaffold theorem, and describe the self-adapting and confidence-tracking extensions.

3.1 The 26 Diagnostic Levels

The holistic diagnostic H is built from 26 levels L_0, L_1, \dots, L_{25} , each measuring a different aspect of the fluid state. Every level is computed as a *dual number* (v, v') where v is the value and $v' = dv/d\theta$ is the derivative with respect to the state parameter θ . This is implemented via forward-mode automatic differentiation, which is native to the Simplex programming language.

Definition 7 (Dual Number). *A dual number is a pair $(v, v') \in \mathbb{R} \times \mathbb{R}$ with arithmetic:*

$$(a, a') + (b, b') = (a + b, a' + b'), \quad (27)$$

$$(a, a') \cdot (b, b') = (ab, ab' + a'b), \quad (28)$$

$$f(a, a') = (f(a), f'(a) \cdot a'). \quad (29)$$

The dual part v' propagates exact first derivatives through arbitrary compositions of differentiable functions without numerical approximation.

We now define each level in detail.

3.1.1 Level 0: State Vector θ

Definition 8 (State Vector). *The state vector is $\theta(t) = (u_1, \dots, u_K, \omega_1, \dots, \omega_K) \in \mathbb{R}^{2K}$, where K is the number of wavenumber modes. As a dual number:*

$$L_0(t) = (\theta(t), \dot{\theta}(t)), \quad (30)$$

where $\dot{\theta}(t)$ is the time derivative computed from the ODE right-hand side (23)–(24).

The derivative $\dot{\theta}$ tells us the instantaneous evolution direction. L_0 is always active (gate always open) and serves as the foundation from which all other levels are computed. Its informative content is the full phase-space trajectory, which is too high-dimensional for direct classification but provides the raw data for all higher levels.

3.1.2 Level 1: Normalised Enstrophy

Definition 9 (Normalised Enstrophy). *The normalised enstrophy is:*

$$L_1(t) = \frac{\Omega(t)}{\Omega(0)}, \quad \Omega(t) = \sum_{i=1}^K \omega_i(t)^2. \quad (31)$$

As a dual number:

$$L_1 = \left(\frac{\Omega}{\Omega(0)}, \frac{2}{\Omega(0)} \sum_i \omega_i \dot{\omega}_i \right). \quad (32)$$

Physical meaning: L_1 measures how much the total vorticity has grown relative to its initial value. $L_1 = 1$ means the enstrophy is at its initial level; $L_1 \gg 1$ indicates significant vorticity amplification; $L_1 > 100$ (our blow-up threshold) indicates that vorticity has grown by two orders of magnitude.

Dual part: The derivative $L_1' = dL_1/dt$ gives the instantaneous enstrophy growth rate. Positive L_1' indicates growing enstrophy; the sign and magnitude of L_1' provide early warning of potential blow-up.

Gate condition: L_1 is always active ($g_1 = 1$ for all t).

Connection to other levels: L_1 is the primary driver of the feedback loop. Rising L_1 increases L_4 (fragility), which degrades L_2 (convergence), which permits further L_1 growth.

3.1.3 Level 2: Convergence Score

Definition 10 (Convergence Score). *The convergence score measures how much the enstrophy trajectory is stabilising over time:*

$$L_2(t) = S(t) = 1 - \frac{D_{\text{late}}(t)}{D_{\text{early}}(t)}, \quad (33)$$

where $D_{\text{early}}(t)$ and $D_{\text{late}}(t)$ are the mean absolute enstrophy drift over the first and second halves of the observation window $[0, t]$:

$$D_{\text{early}}(t) = \frac{2}{t} \int_0^{t/2} |\dot{\Omega}(s)| ds, \quad D_{\text{late}}(t) = \frac{2}{t} \int_{t/2}^t |\dot{\Omega}(s)| ds. \quad (34)$$

Physical meaning: $S \rightarrow 1$ means late drift is much smaller than early drift — the solution is converging to a steady state. $S \rightarrow 0$ means late drift equals early drift — the solution shows no convergence. $S < 0$ means late drift exceeds early drift — the solution is *diverging*.

Derivation: The ratio $D_{\text{late}}/D_{\text{early}}$ measures the relative “activity” of the trajectory in its second half compared to its first half. For a solution approaching equilibrium, the activity should decrease with time, giving $S > 0$. For a solution approaching blow-up, the activity increases, giving $S < 0$.

Dual part: The derivative $S'(t) = dS/dt$ indicates whether convergence is improving ($S' > 0$) or degrading ($S' < 0$). A positive S' is a healthy sign; a negative S' signals potential trouble even if S itself is currently positive.

Gate condition: $g_2(t) = 1$ when $t > t_{\min}$ (a minimum observation window is required for the early/late comparison to be meaningful). We use $t_{\min} = 1,000$ steps.

3.1.4 Level 3: PID Control on Convergence

Definition 11 (Convergence Derivatives). L_3 provides PID (Proportional-Integral-Derivative) control signals on the convergence score S :

$$L_3^{(P)}(t) = S(t), \quad (35)$$

$$L_3^{(I)}(t) = \frac{1}{t} \int_0^t S(\tau) d\tau, \quad (36)$$

$$L_3^{(D)}(t) = S'(t) = \frac{dS}{dt}. \quad (37)$$

Physical meaning: The PID decomposition separates three timescales of information about convergence:

- $L_3^{(P)}$: Current convergence state (instantaneous).
- $L_3^{(I)}$: Historical convergence (integrated over the trajectory). If $L_3^{(I)} < 0$ despite $L_3^{(P)} > 0$, the trajectory has spent more time diverging than converging — a warning sign.
- $L_3^{(D)}$: Convergence rate of change (predictive). A sharply negative $L_3^{(D)}$ predicts imminent convergence collapse even if $L_3^{(P)}$ is currently positive.

Dual part: Each component is itself a dual number. The second derivative S'' (the derivative of the derivative) measures the *acceleration* of convergence change, providing a third-order predictive signal.

Gate condition: $g_3(t) = 1$ when $t > 2t_{\min}$.

3.1.5 Level 4: Stability Margin

Definition 12 (Stability Margin). *The stability margin measures the sensitivity of the trajectory to perturbation:*

$$L_4(t) = M(\theta, \varepsilon) = \frac{\|\theta_\varepsilon(t) - \theta(t)\|}{\varepsilon}, \quad (38)$$

where $\theta(t)$ is the unperturbed trajectory and $\theta_\varepsilon(t)$ is the trajectory starting from $\theta(0) + \varepsilon\hat{e}$ for a random unit perturbation direction \hat{e} .

Physical meaning: M is the amplification factor of a small perturbation over the time interval $[0, t]$. $M \approx 1$ means the system is neutrally stable (perturbations neither grow nor shrink). $M \gg 1$ means the system is highly sensitive to initial conditions — it is chaotic or approaching a singularity. $M < 1$ means the system is contracting (perturbations are damped).

Derivation: M is computed by running two simulations in parallel: the original trajectory $\theta(t)$ and the perturbed trajectory $\theta_\varepsilon(t) = \theta(t) + \delta(t)$, where $\delta(0) = \varepsilon\hat{e}$. The perturbation $\delta(t)$ evolves according to the linearised dynamics:

$$\dot{\delta} = J(\theta(t)) \cdot \delta + O(\delta^2), \quad (39)$$

where $J(\theta) = \partial f / \partial \theta$ is the Jacobian of the ODE right-hand side. For small ε , the ratio $\|\delta(t)\| / \varepsilon$ converges to the maximum singular value of the state transition matrix $\Phi(t, 0)$.

Dual part: The derivative $M' = dM/dt$ gives the instantaneous growth rate of fragility. $M' > 0$ indicates growing instability; $M' < 0$ indicates stabilisation.

Gate condition: $g_4(t) = 1$ for $t > 0$ (always active after the perturbation is applied at $t = 0$).

Connection to other levels: M is driven by L_1 (higher enstrophy \rightarrow larger Jacobian eigenvalues \rightarrow larger M) and drives L_2 (larger $M \rightarrow$ less predictable trajectory \rightarrow lower convergence score).

3.1.6 Level 5: Fragility Evolution

Definition 13 (Fragility Rate).

$$L_5(t) = \frac{dM}{dt}(\theta, \varepsilon, t). \quad (40)$$

Physical meaning: L_5 measures how rapidly fragility is changing. A positive L_5 indicates growing instability; a negative L_5 indicates that the system is becoming more robust. The key diagnostic use is detecting the *onset* of instability: L_5 transitions from negative to positive before M itself becomes large, providing early warning.

Dual part: $L_5' = d^2M/dt^2$ measures the acceleration of fragility growth. Positive L_5' indicates that fragility is growing faster over time — a hallmark of superexponential instability.

Gate condition: $g_5(t) = g_4(t)$ (same as L_4).

3.1.7 Level 6: Perturbation Scaling

Definition 14 (Perturbation Scaling).

$$L_6(t) = \frac{dM}{d\varepsilon}(t) = \lim_{\delta\varepsilon \rightarrow 0} \frac{M(\theta, \varepsilon + \delta\varepsilon, t) - M(\theta, \varepsilon, t)}{\delta\varepsilon}. \quad (41)$$

Physical meaning: L_6 measures whether the perturbation response is linear ($L_6 \approx 0$, meaning M is independent of ε) or superlinear ($L_6 > 0$, meaning larger perturbations are amplified disproportionately). Superlinear scaling indicates that the system is approaching a nonlinear instability, where small perturbations behave differently from large ones.

Computation: In practice, we compute L_6 by finite differences using three perturbation amplitudes $\varepsilon/2$, ε , and 2ε :

$$L_6 \approx \frac{M(\theta, 2\varepsilon) - M(\theta, \varepsilon/2)}{3\varepsilon/2}. \quad (42)$$

Dual part: $L'_6 = d^2M/(dt d\varepsilon)$ is the mixed second derivative, measuring how the nonlinearity of the perturbation response changes with time.

Gate condition: $g_6(t) = 1$ when $M > 2$ (the perturbation response is large enough for the scaling analysis to be meaningful).

3.1.8 Level 7: Optimal Probe

Definition 15 (Optimal Probe).

$$L_7 = \varepsilon^* = \arg \min_{\varepsilon} \text{Var}[M(\theta, \varepsilon)] \quad (43)$$

where the variance is over multiple random perturbation directions \hat{e} .

Physical meaning: ε^* is the perturbation amplitude that gives the most consistent (least noisy) stability margin measurement. Too small, and numerical noise dominates; too large, and nonlinear effects make the measurement direction-dependent. The optimal probe is the “right question to ask” about the system’s stability.

Implementation: ε^* is learned via a meta-gradient:

$$\varepsilon_{n+1}^* = \varepsilon_n^* - \eta \frac{\partial}{\partial \varepsilon} \text{Var}[M(\theta, \varepsilon)] \Big|_{\varepsilon = \varepsilon_n^*}. \quad (44)$$

The gradient is computed by forward-mode AD through the variance computation. In practice, ε^* converges within 5–10 iterations and remains stable thereafter.

Gate condition: $g_7(t) = 1$ when $t > 5,000$ steps (sufficient trajectory for variance estimation).

3.1.9 Level 8: Simulation Parameter Sensitivity

Definition 16 (Parameter Sensitivity).

$$L_8 = \left(\frac{\partial T_{\text{blow}}}{\partial A}, \frac{\partial T_{\text{blow}}}{\partial \nu} \right), \quad (45)$$

where T_{blow} is the blow-up time (or T if no blow-up occurs).

Physical meaning: $\partial T_{\text{blow}}/\partial A$ measures how sensitively the blow-up time depends on the initial amplitude. Near A^* , this derivative diverges — small changes in amplitude cause large changes in blow-up time. Far from A^* (either below or above), the derivative is moderate. The sensitivity $\partial T_{\text{blow}}/\partial \nu$ similarly measures viscosity sensitivity.

Computation: These derivatives are computed by finite differences:

$$\frac{\partial T_{\text{blow}}}{\partial A} \approx \frac{T_{\text{blow}}(A + \delta A) - T_{\text{blow}}(A - \delta A)}{2\delta A}, \quad (46)$$

$$\frac{\partial T_{\text{blow}}}{\partial \nu} \approx \frac{T_{\text{blow}}(\nu + \delta \nu) - T_{\text{blow}}(\nu - \delta \nu)}{2\delta \nu}, \quad (47)$$

with $\delta A = 10^{-4}$ and $\delta \nu = 10^{-5}$.

Gate condition: $g_8(t) = 1$ at $t = T$ (end of simulation only).

3.1.10 Level 9: Physics Coupling Sensitivity

Definition 17 (Coupling Sensitivity).

$$L_9 = \left(\frac{\partial T_{\text{blow}}}{\partial \lambda}, \frac{\partial^2 T_{\text{blow}}}{\partial \lambda \partial A} \right), \quad (48)$$

where λ is the linear stretching coefficient.

Physical meaning: The first component measures how the blow-up time depends on the stretching strength. The cross-derivative $\partial^2 T / (\partial \lambda \partial A)$ measures how the amplitude sensitivity changes with stretching strength — a second-order interaction effect. Large cross-derivatives indicate that the system’s response to amplitude changes depends strongly on the physics coupling, i.e., the stretching and amplitude are not independent factors.

Computation: The cross-derivative is computed by a 2×2 finite difference stencil:

$$\frac{\partial^2 T}{\partial \lambda \partial A} \approx \frac{T(\lambda + \delta \lambda, A + \delta A) - T(\lambda + \delta \lambda, A - \delta A) - T(\lambda - \delta \lambda, A + \delta A) + T(\lambda - \delta \lambda, A - \delta A)}{4\delta \lambda \delta A}. \quad (49)$$

Gate condition: $g_9(t) = 1$ at $t = T$.

3.1.11 Level 10: Trajectory Acceleration

Definition 18 (Trajectory Acceleration).

$$L_{10}(t) = \frac{\partial T_{\text{blow}}}{\partial \alpha_\nu}, \quad (50)$$

where α_ν is the rate parameter in a viscosity sweep $\nu(t) = \nu_0 + \alpha_\nu \cdot t$.

Physical meaning: L_{10} measures how the trajectory responds to time-varying viscosity — a “chirp” sweep through the viscosity spectrum. Positive L_{10} indicates that increasing viscosity over time delays blow-up (as expected); the magnitude of L_{10} indicates how responsive the system is to viscosity changes. The acceleration perspective reveals the system’s *adaptability*: a system with large L_{10} can be “steered” away from blow-up by viscosity adjustment, while a system with small L_{10} is committed to its fate regardless of viscosity changes.

Dual part: $L'_{10} = \partial^2 T / (\partial \alpha_\nu \partial t)$ measures how the viscosity sensitivity changes over time, identifying windows where viscosity intervention is most effective.

Gate condition: $g_{10}(t) = 1$ when viscosity sweep is active.

3.1.12 Level 11: Enstrophy Acceleration

Definition 19 (Enstrophy Acceleration).

$$L_{11}(t) = \frac{d^2 \Omega}{dt^2} = 2 \sum_i (\dot{\omega}_i^2 + \omega_i \ddot{\omega}_i), \quad (51)$$

where $\ddot{\omega}_i$ is computed by differentiating the ODE right-hand side with respect to time (using the chain rule through the dual-number arithmetic).

Physical meaning: L_{11} measures whether enstrophy growth is accelerating ($L_{11} > 0$) or decelerating ($L_{11} < 0$). Even if enstrophy is growing ($d\Omega/dt > 0$), deceleration ($L_{11} < 0$) indicates that the growth is self-limiting and will eventually stop. Conversely, positive acceleration indicates that the growth rate is itself increasing — a precursor to runaway.

Dual part: $L'_{11} = d^3\Omega/dt^3$ is the enstrophy jerk (Level 13), providing one further level of prediction.

Gate condition: $g_{11}(t) = 1$ when $d\Omega/dt > 0$ (only informative when enstrophy is actually growing).

3.1.13 Level 12: Deceleration Ratio

Definition 20 (Deceleration Ratio).

$$L_{12}(t) = \frac{d^2\Omega/dt^2}{(d\Omega/dt)^2}, \quad (52)$$

with the convention $L_{12} = 0$ when $d\Omega/dt = 0$.

Physical meaning: L_{12} normalises the acceleration by the square of the growth rate. For exponential growth $\Omega(t) = \Omega_0 e^{\gamma t}$, we have $L_{12} = 1/\Omega_0 e^{\gamma t} \rightarrow 0$ as $t \rightarrow \infty$. For power-law growth $\Omega(t) = C(T-t)^{-\alpha}$, $L_{12} = (\alpha+1)/(\alpha\Omega)$. The ratio distinguishes between self-similar blow-up (constant L_{12}) and exponential growth (decaying L_{12}), which have very different implications for regularity.

Gate condition: $g_{12}(t) = g_{11}(t)$ (same as L_{11}).

3.1.14 Level 13: Enstrophy Jerk

Definition 21 (Enstrophy Jerk).

$$L_{13}(t) = \frac{d^3\Omega}{dt^3}. \quad (53)$$

Physical meaning: The third derivative of enstrophy. Positive jerk ($L_{13} > 0$) combined with positive acceleration ($L_{11} > 0$) and positive growth ($d\Omega/dt > 0$) creates a triple-positive signature that is a strong indicator of imminent blow-up. Safe trajectories typically have negative jerk (the growth rate is decelerating at an increasing rate), while blow-up trajectories have positive jerk (the growth rate is accelerating at an increasing rate).

Computation: Computed by forward-mode AD through three compositions of the ODE right-hand side, using the dual-number tower (v, v', v'', v''') .

Gate condition: $g_{13}(t) = g_{11}(t)$.

3.1.15 Level 14: Saturation Time Predictor

Definition 22 (Saturation Time).

$$L_{14}(t) = t_{sat}(t) = -\frac{d\Omega/dt}{d^2\Omega/dt^2}, \quad (54)$$

defined when $d^2\Omega/dt^2 < 0$ (i.e., when growth is decelerating).

Derivation: t_{sat} estimates the time at which enstrophy growth will reach zero, assuming the current deceleration continues linearly. If $d\Omega/dt = v$ and $d^2\Omega/dt^2 = a < 0$, then the growth rate $v(t) \approx v_0 + at$ reaches zero at $t = -v_0/a = t_{sat}$. This is a first-order prediction; higher-order predictions use L_{13} (jerk) to account for nonlinear deceleration.

Physical meaning: t_{sat} is the predicted time to enstrophy saturation. A finite, moderate t_{sat} indicates that enstrophy growth is self-limiting; $t_{sat} \rightarrow \infty$ (or undefined, because $d^2\Omega/dt^2 \geq 0$) indicates that growth shows no sign of stopping.

Dual part: $L'_{14} = dt_{\text{sat}}/dt$ measures whether the saturation prediction is getting closer ($L'_{14} < 0$) or receding ($L'_{14} > 0$). A receding saturation time is a warning sign.

Gate condition: $g_{14}(t) = 1$ when $d^2\Omega/dt^2 < 0$ (only defined during deceleration).

3.1.16 Level 15: Enstrophy Doubling Time

Definition 23 (Doubling Time). *The n -th enstrophy doubling time is:*

$$\tau_d(n) = t_n - t_{n-1}, \quad \text{where } \Omega(t_n) = 2^n \cdot \Omega(0). \quad (55)$$

L_{15} is the pair $(\tau_d(n), \tau_d(n)/\tau_d(n-1))$.

Physical meaning: τ_d is the time for enstrophy to double. For exponential growth $\Omega = \Omega_0 e^{\gamma t}$, $\tau_d = \ln 2/\gamma$ is constant. For superexponential growth, τ_d shrinks with each doubling. For subexponential (power-law) growth, τ_d grows with each doubling.

The ratio $\tau_d(n)/\tau_d(n-1)$ is the key diagnostic:

- $\tau_d(n)/\tau_d(n-1) > 1$: decelerating growth (safe).
- $\tau_d(n)/\tau_d(n-1) = 1$: exponential growth (marginal).
- $\tau_d(n)/\tau_d(n-1) < 0.85$: accelerating growth (blow-up).

Dual part: The derivative $d\tau_d/dn$ (treating n as continuous) measures the trend in doubling time. This is approximated by $\tau_d(n) - \tau_d(n-1)$.

Gate condition: $g_{15}(t) = 1$ when at least two doublings have occurred.

This level is developed in full detail in Section 4.

3.1.17 Level 16: Reynolds Fluctuation Ratio

Definition 24 (Reynolds Fluctuation Ratio).

$$L_{16}(t) = \frac{|\Omega(t) - \langle \Omega \rangle_W|}{\langle \Omega \rangle_W}, \quad (56)$$

where $\langle \Omega \rangle_W = \frac{1}{W} \int_{t-W}^t \Omega(s) ds$ is the windowed time average over W steps.

Physical meaning: This is inspired by RANS turbulence modelling, where the Reynolds decomposition separates the flow into mean and fluctuating parts: $\Omega = \bar{\Omega} + \Omega'$. The ratio $|\Omega'|/\bar{\Omega}$ measures the relative magnitude of the fluctuations. In a steady turbulent flow, this ratio is $O(1)$. In a solution approaching blow-up, the fluctuations grow faster than the mean, and the ratio increases without bound.

Dual part: $L'_{16} = d/dt(|\Omega - \bar{\Omega}|/\bar{\Omega})$ measures whether the fluctuation ratio is growing (indicating increasing unsteadiness) or shrinking (indicating approach to a statistically steady state).

Gate condition: $g_{16}(t) = 1$ when $t > W$ (sufficient data for time averaging). We use $W = 5,000$ steps.

3.1.18 Level 17: Turbulent Timescale Ratio

Definition 25 (Timescale Ratio).

$$L_{17}(t) = \frac{k(t)/\varepsilon(t)}{\tau_d(t)}, \quad (57)$$

where $k = \frac{1}{2} \sum_i u_i^2$ is the kinetic energy and $\varepsilon = \nu \sum_i k_i^2 u_i^2$ is the dissipation rate.

Physical meaning: k/ε is the turbulent timescale from the k - ε model [11]. τ_d is the enstrophy doubling time. The ratio L_{17} compares these two timescales:

- $L_{17} \gg 1$: turbulent timescale is much longer than doubling time, meaning enstrophy grows faster than the turbulent cascade can redistribute energy. This is a blow-up precursor.
 - $L_{17} \sim 1$: timescales are comparable, indicating equilibrium between enstrophy production and turbulent dissipation.
 - $L_{17} \ll 1$: doubling time is much longer than the turbulent timescale, meaning the cascade equilibrates fast and enstrophy growth is slow. This is the safe regime.
- Gate condition:** $g_{17}(t) = g_{15}(t)$ (requires a well-defined doubling time).

3.1.19 Level 18: Relaxation Deficit

Definition 26 (Relaxation Deficit).

$$L_{18}(t) = \frac{|F_{stretch}(t)|}{|F_{diff}(t)|} = \frac{\sum_i |F_{stretch}^{(i)}|}{\sum_i \nu k_i^2 |\omega_i|}, \quad (58)$$

where $F_{stretch}^{(i)}$ is the total stretching force on mode i and $F_{diff}^{(i)} = -\nu k_i^2 \omega_i$ is the diffusion force.

Physical meaning: Inspired by the Lattice Boltzmann method, where the collision operator drives the distribution function toward local equilibrium. The relaxation deficit measures how far the system is from the “equilibrium” where stretching and diffusion balance. $L_{18} < 1$ means diffusion dominates (the system is over-relaxed); $L_{18} > 1$ means stretching dominates (the system is under-relaxed, accumulating enstrophy). The critical ratio $L_{18} = 1$ corresponds to exact balance.

Dual part: $L'_{18} = d/dt(|F_{stretch}|/|F_{diff}|)$ measures whether the system is moving toward equilibrium ($L'_{18} < 0$ when $L_{18} > 1$) or away from it.

Gate condition: $g_{18}(t) = 1$ for all $t > 0$.

3.1.20 Level 19: Symmetry Breaking

Definition 27 (Symmetry Breaking).

$$L_{19}(t) = \frac{|\omega_x - \omega_y| + |\omega_y - \omega_z| + |\omega_z - \omega_x|}{|\omega_x| + |\omega_y| + |\omega_z|}, \quad (59)$$

where $\omega_x, \omega_y, \omega_z$ are the vorticity components at the lowest wavenumber ($k = 1$).

Physical meaning: In the initial condition, all three vorticity components have similar magnitudes (the initial state is approximately isotropic). As the solution evolves, vortex stretching preferentially amplifies certain directions, breaking the isotropy. L_{19} measures the degree of this symmetry breaking:

- $L_{19} = 0$: perfect isotropy ($\omega_x = \omega_y = \omega_z$).
- $L_{19} \rightarrow 2$: extreme anisotropy (one component dominates).

Strong symmetry breaking is associated with the formation of coherent vortex structures (tubes, sheets) that are the precursors of blow-up.

Gate condition: $g_{19}(t) = 1$ for all $t > 0$.

3.1.21 Level 20: Directional Alignment

Definition 28 (Directional Alignment).

$$L_{20}(t) = \frac{|\dot{\omega} \cdot \omega|}{|\dot{\omega}| \cdot |\omega|}, \quad (60)$$

where $\dot{\omega} = (\dot{\omega}_1, \dots, \dot{\omega}_K)$ is the vorticity rate of change and the dot product and norms are in \mathbb{R}^K .

Physical meaning: L_{20} is the cosine of the angle between the vorticity vector ω and its time derivative $\dot{\omega}$. $L_{20} \approx 1$ means $\dot{\omega}$ is parallel to ω — the vorticity is growing in the same direction it already points. This is “dimension collapse”: the dynamics are effectively one-dimensional, with all growth concentrated in a single direction. This is the signature of vortex tube stretching, where a cylindrical vortex is stretched along its axis.

$L_{20} \approx 0$ means $\dot{\omega}$ is perpendicular to ω — the vorticity vector is rotating without growing. This is vortex tilting, which redistributes vorticity among directions without net amplification.

Gate condition: $g_{20}(t) = 1$ when $|\omega| > 0.1 \cdot \Omega(0)^{1/2}$.

3.1.22 Level 21: Stretching Efficiency

Definition 29 (Stretching Efficiency).

$$L_{21}(t) = \frac{\omega^T S \omega}{|\omega|^2 \cdot |S|}, \quad (61)$$

where $S_{ij} = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$ is the strain rate tensor and $|S| = (S_{ij} S_{ij})^{1/2}$ is the Frobenius norm.

Physical meaning: The enstrophy production term in (4) is $\int \omega_i \omega_j S_{ij} dx = \omega^T S \omega$ in the Galerkin truncation. The stretching efficiency normalises this by $|\omega|^2 |S|$ to obtain the alignment between vorticity and strain. By the Cauchy–Schwarz inequality, $L_{21} \leq 1$, with equality when the vorticity vector is aligned with the principal extensional strain direction.

In real turbulence, DNS data shows that $L_{21} \approx 0.1$ – 0.2 on average [14], with intermittent excursions to higher values near coherent structures. In our models, L_{21} is the *earliest discriminator* between safe and blow-up trajectories: it begins to diverge at step $\sim 6,000$, well before enstrophy growth becomes apparent.

Computation in the Galerkin model: The strain rate tensor is approximated from the velocity modes. For the truncated system, we construct S from the coupling coefficients:

$$S_{ij}^{(\text{model})} = \frac{1}{2}(c_{ij} u_j + c_{ji} u_i). \quad (62)$$

Gate condition: $g_{21}(t) = 1$ for all $t > 0$.

3.1.23 Level 22: Back-Reaction

Definition 30 (Back-Reaction).

$$L_{22}(t) = \frac{\sum_i |u_i(t) - u_i^{(\text{passive})}(t)|}{\sum_i |u_i(t)|}, \quad (63)$$

where $u_i^{(\text{passive})}$ is the velocity evolved without the vorticity feedback (i.e., setting $\sigma = 0$).

Physical meaning: L_{22} measures how much the vorticity dynamics have modified the velocity field compared to passive evolution. In a regular solution, the back-reaction is bounded; in a blow-up scenario, the back-reaction grows without limit because the velocity field is increasingly driven by the diverging vorticity.

Gate condition: $g_{22}(t) = 1$ for $t > 1,000$.

3.1.24 Level 23: Energy Transfer Flux

Definition 31 (Energy Transfer Flux).

$$L_{23}(t) = \Pi(k_c, t) = \sum_{k_i > k_c} \sum_{k_j \leq k_c} \beta_{ji} \omega_j^2 \cdot \text{sgn}(\omega_i) \cdot \omega_i, \quad (64)$$

where $k_c = K/2$ is a cutoff wavenumber separating large and small scales.

Physical meaning: $\Pi(k_c)$ is the energy flux across the wavenumber k_c , representing the rate of energy transfer from scales larger than $1/k_c$ to scales smaller than $1/k_c$. In the Kolmogorov inertial range, Π is constant and equals the dissipation rate ε . In our models, Π provides a direct measure of the forward cascade intensity:

- $\Pi > 0$: net forward cascade (large \rightarrow small), normal turbulent behaviour.
- $\Pi \gg \varepsilon$: cascade bottleneck, enstrophy accumulating at small scales faster than it can be dissipated.
- $\Pi < 0$: inverse cascade (small \rightarrow large), rare in 3D.

Gate condition: $g_{23}(t) = 1$ when $K \geq 4$ (at least 4 wavenumbers needed for meaningful scale separation).

3.1.25 Level 24: Phase Coherence

Definition 32 (Phase Coherence).

$$L_{24}(t) = \frac{|\sum_i \omega_i(t) e^{ik_i \phi(t)}|}{\sum_i |\omega_i(t)|}, \quad (65)$$

where $\phi(t) = \arg(\omega_1(t) + i\omega_2(t))$ is a reference phase.

Physical meaning: Phase coherence measures the degree to which the vorticity modes are phase-locked. $L_{24} \approx 1$ indicates strong coherence (all modes oscillate in phase), associated with coherent vortex structures. $L_{24} \approx 0$ indicates random phases (incoherent turbulence). Phase coherence is a precursor to blow-up: coherent structures amplify vortex stretching because the stretching term $\omega_i \omega_j S_{ij}$ is maximised when the modes are in phase.

Gate condition: $g_{24}(t) = 1$ when $K \geq 3$.

3.1.26 Level 25: Spectral Slope

Definition 33 (Spectral Slope).

$$L_{25}(t) = - \left. \frac{d \ln E(k)}{d \ln k} \right|_{fit}, \quad (66)$$

where $E(k_i) = \frac{1}{2}(u_i^2 + \omega_i^2/k_i^2)$ is the energy spectrum and the slope is obtained by least-squares fitting $\ln E$ vs $\ln k$.

Physical meaning: The spectral slope characterises the energy distribution across scales. In the Kolmogorov inertial range, the slope is $5/3 \approx 1.67$. Steeper slopes indicate faster energy decay at small scales (dominated by dissipation); shallower slopes indicate energy accumulation at small scales (precursor to blow-up). A slope that decreases over time indicates that energy is building up at high wavenumbers, consistent with the enstrophy cascade.

Gate condition: $g_{25}(t) = 1$ when $K \geq 4$.

3.1.27 Summary of Levels

3.2 The Feedback Loop

3.2.1 Enstrophy Production Equation

The enstrophy production equation is the starting point for understanding the feedback loop. From the ODE system (24), the enstrophy $\Omega = \sum_i \omega_i^2$ evolves as:

$$\frac{d\Omega}{dt} = 2 \sum_i \omega_i \dot{\omega}_i = \underbrace{-2\nu \sum_i k_i^2 \omega_i^2}_{D(t)} + \underbrace{2\lambda \sum_i \omega_i \sum_j c_{ij} \omega_j u_j}_{P_1(t)} + \underbrace{2\lambda_2 |\omega|^2 \sum_i \omega_i^2}_{P_2(t)} + \underbrace{2 \sum_i \omega_i F_{\text{cascade}}^{(i)}}_{C(t)}. \quad (67)$$

Table 1: Summary of the 26 diagnostic levels.

Level	Name	Equation	Dual Part Meaning	Gate
L_0	State vector	$\theta = (u, \omega)$	Velocity in phase space	Always
L_1	Enstrophy	$\Omega/\Omega(0)$	Growth rate	Always
L_2	Convergence	$1 - D_{\text{late}}/D_{\text{early}}$	Convergence trend	$t > t_{\text{min}}$
L_3	PID control	(S, \bar{S}, S')	Convergence dynamics	$t > 2t_{\text{min}}$
L_4	Stability margin	$\ \delta\theta\ /\varepsilon$	Fragility rate	Always
L_5	Fragility rate	dM/dt	Fragility acceleration	Always
L_6	Perturbation scaling	$dM/d\varepsilon$	Nonlinearity trend	$M > 2$
L_7	Optimal probe	ε^*	Probe adaptation rate	$t > 5000$
L_8	Parameter sensitivity	$\partial T/\partial(A, \nu)$	Sensitivity trend	$t = T$
L_9	Coupling sensitivity	$\partial T/\partial\lambda$, cross	Interaction dynamics	$t = T$
L_{10}	Trajectory accel.	$\partial T/\partial\alpha_\nu$	Sweep effectiveness	Sweep active
L_{11}	Enstrophy accel.	$d^2\Omega/dt^2$	Jerk	$\dot{\Omega} > 0$
L_{12}	Deceleration ratio	$\ddot{\Omega}/\dot{\Omega}^2$	Blow-up type	$\dot{\Omega} > 0$
L_{13}	Jerk	$d^3\Omega/dt^3$	Snap	$\dot{\Omega} > 0$
L_{14}	Saturation time	$-\dot{\Omega}/\ddot{\Omega}$	Predicted settling	$\ddot{\Omega} < 0$
L_{15}	Doubling time	τ_d and ratio	Trend	2+ doublings
L_{16}	Reynolds ratio	$ \Omega - \bar{\Omega} /\bar{\Omega}$	Fluctuation trend	$t > W$
L_{17}	Timescale ratio	$(k/\varepsilon)/\tau_d$	Equilibration	τ_d defined
L_{18}	Relaxation deficit	$ F_{\text{str}} / F_{\text{diff}} $	Balance trend	Always
L_{19}	Symmetry breaking	$\sum \omega_i - \omega_j /\sum \omega_i $	Anisotropy growth	Always
L_{20}	Alignment	$ \dot{\omega} \cdot \omega /(\dot{\omega} \omega)$	Collapse trend	$ \omega $ large
L_{21}	Stretching eff.	$\omega^T S \omega / (\omega ^2 S)$	Alignment trend	Always
L_{22}	Back-reaction	$ u - u^{\text{pass}} / u $	Growth trend	$t > 1000$
L_{23}	Energy flux	$\Pi(k_c)$	Flux trend	$K \geq 4$
L_{24}	Phase coherence	$ \sum \omega_i e^{ik_i \phi} /\sum \omega_i $	Coherence trend	$K \geq 3$
L_{25}	Spectral slope	$-d \ln E/d \ln k$	Slope evolution	$K \geq 4$

Here $D(t) < 0$ is the diffusion (always negative, always dissipating enstrophy), $P_1(t)$ is the linear stretching production, $P_2(t) = 2\lambda_2\Omega^2 > 0$ is the quadratic production (always positive when $\Omega > 0$), and $C(t)$ is the cascade contribution.

The critical observation is that $P_2(t) = 2\lambda_2\Omega^2$ grows as the *square* of enstrophy, while $D(t)$ grows only linearly in enstrophy (since $\sum k_i^2\omega_i^2 \leq k_K^2\Omega$). This creates the possibility of runaway: once Ω is large enough that $P_2 > |D|$, the production exceeds diffusion and enstrophy grows, which makes P_2 even larger, creating a positive feedback loop.

3.2.2 How L_1 Drives L_4

Lemma 34 (Enstrophy–Fragility Coupling). *If enstrophy $\Omega(t)$ is growing, the stability margin $M(t)$ increases at a rate bounded below by:*

$$\frac{dM}{dt} \geq \frac{\lambda_2\Omega}{k_K^2\nu} \cdot M, \quad (68)$$

when Ω is sufficiently large.

Proof. The perturbation $\delta\theta$ evolves according to the linearised dynamics (39), with Jacobian $J(\theta)$. The largest eigenvalue of J is dominated by the quadratic term $\lambda_2|\omega|^2$, which contributes $2\lambda_2\Omega$ to the diagonal of J . The diffusion contributes $-\nu k_K^2$ to the diagonal. For Ω large enough that $2\lambda_2\Omega > \nu k_K^2$ (the maximum diffusion), the maximum eigenvalue of J is at least $2\lambda_2\Omega - \nu k_K^2 > \lambda_2\Omega$ (for $\Omega > \nu k_K^2/\lambda_2$). The stability margin $M(t) \sim e^{\int_0^t \lambda_{\max}(J) ds}$, giving $dM/dt \geq \lambda_{\max}(J) \cdot M \geq \lambda_2\Omega \cdot M / (k_K^2\nu)$ after normalising. \square

The physical content of Lemma 34 is that high enstrophy makes the system fragile: small perturbations are amplified by the nonlinear stretching term, and the amplification rate grows with Ω .

3.2.3 How L_4 Drives L_2

Lemma 35 (Fragility–Convergence Coupling). *If the stability margin $M(t)$ grows exponentially ($M(t) \geq M_0 e^{\gamma t}$ for some $\gamma > 0$), then the convergence score $S(t)$ degrades:*

$$S(t) \leq 1 - \frac{M(t)}{M(t/2)} \cdot \frac{D_{\text{early}}(t/2)}{D_{\text{early}}(t)}. \quad (69)$$

Proof. When M grows exponentially, the trajectory becomes increasingly sensitive to perturbations, which means that the enstrophy trajectory also becomes more variable. Specifically, the late drift D_{late} is amplified by the fragility: different perturbation realisations lead to different late-time enstrophy values, with spread proportional to $M(t) \cdot \varepsilon$. Since D_{late} measures the absolute drift of enstrophy, and this drift scales with the perturbation amplification:

$$D_{\text{late}}(t) \geq C \cdot M(t) \cdot |d\Omega/d\theta|,$$

while $D_{\text{early}}(t)$ is bounded by the initial perturbation response. The convergence score $S = 1 - D_{\text{late}}/D_{\text{early}}$ therefore decreases as M grows. \square

3.2.4 How L_2 Drives L_1

Lemma 36 (Convergence–Enstrophy Coupling). *When the convergence score $S(t) < 0$ (the trajectory is diverging), the enstrophy growth rate satisfies:*

$$\frac{d\Omega}{dt} \geq (1 - S(t)) \cdot \left. \frac{d\Omega}{dt} \right|_{\text{baseline}}, \quad (70)$$

where the baseline growth rate is the rate at $S = 0$.

Proof. When $S < 0$, the late drift exceeds the early drift, meaning the enstrophy changes are *accelerating*. The factor $(1 - S)$ quantifies this acceleration: $S = -1$ means the late drift is twice the early drift, amplifying the growth rate by a factor of 2.

More precisely, the convergence score affects enstrophy through the cascade mechanism. When the trajectory is converging ($S > 0$), the stretching and diffusion terms are approaching balance, which constrains enstrophy growth. When convergence fails ($S < 0$), the imbalance grows, and the production term $P_2 = 2\lambda_2\Omega^2$ increasingly exceeds the diffusion term D , leading to amplified enstrophy growth. \square

3.2.5 Closing the Loop

Theorem 37 (Coupling Topology). *The levels L_1 (enstrophy), L_2 (convergence), and L_4 (fragility) form a positive feedback loop:*

$$L_1 \uparrow \rightarrow L_4 \uparrow \rightarrow L_2 \downarrow \rightarrow L_1 \uparrow$$

This loop is present at every Galerkin truncation from 6 to 24 modes. When the loop gain exceeds 1, the system enters runaway and enstrophy diverges.

Proof. By Lemmas 34, 35, and 36, each link in the loop is established:

1. $L_1 \uparrow \Rightarrow L_4 \uparrow$ (Lemma 34): growing enstrophy increases the Jacobian's maximum eigenvalue, amplifying perturbations.
2. $L_4 \uparrow \Rightarrow L_2 \downarrow$ (Lemma 35): growing fragility makes late-time drift exceed early-time drift, degrading the convergence score.
3. $L_2 \downarrow \Rightarrow L_1 \uparrow$ (Lemma 36): convergence failure permits continued enstrophy growth by preventing the stretching-diffusion balance.

The loop gain G is the product of the three coupling strengths:

$$G = \frac{dL_4}{dL_1} \cdot \frac{dL_2}{dL_4} \cdot \frac{dL_1}{dL_2}. \quad (71)$$

When $G > 1$, a perturbation that increases L_1 is amplified through the loop and returns as an even larger increase in L_1 , creating runaway growth. When $G < 1$, perturbations are attenuated and the system returns to equilibrium.

Computational verification confirms the loop at every mode count from 6 to 24. The coupling matrix, measured at the onset of the growth phase ($\Omega/\Omega(0) \approx 2$), is:

$$\begin{pmatrix} dL'_1/L_1 \\ dL'_4/L_4 \\ dL'_2/L_2 \end{pmatrix} = \begin{pmatrix} 0 & 0 & -0.73 \\ 1.85 & 0 & 0 \\ 0 & -0.62 & 0 \end{pmatrix} \begin{pmatrix} L_1 \\ L_4 \\ L_2 \end{pmatrix}. \quad (72)$$

The eigenvalues of the loop matrix determine stability. For the measured coupling strengths, the product $1.85 \times 0.62 \times 0.73 = 0.84$ gives a loop gain less than 1 at moderate amplitudes ($A \lesssim A^*$), and the coupling strengths increase with amplitude, crossing $G = 1$ at $A = A^*$. \square

Remark 38. The coupling matrix (72) is measured from the 8-mode model at $A = 0.25$ (below $A^* = 0.290$). At $A = 0.30$ (above A^*), the coupling strengths increase to approximately 2.4, 0.78, 0.85, giving loop gain $2.4 \times 0.78 \times 0.85 = 1.59 > 1$. The transition from $G < 1$ to $G > 1$ is sharp and occurs at $A = A^*$ to within the binary search precision of 10^{-5} .

3.2.6 The Loop Gain Computation

The loop gain $G(A)$ as a function of amplitude is computed as follows. For each amplitude A , we run the simulation and measure the three coupling derivatives at the enstrophy doubling point ($\Omega = 2\Omega(0)$, if reached, otherwise at $t = T$):

$$G(A) = \left(\frac{\partial L_4}{\partial L_1} \right) \cdot \left(\frac{\partial L_2}{\partial L_4} \right) \cdot \left(\frac{\partial L_1}{\partial L_2} \right). \quad (73)$$

Each partial derivative is computed by the dual-number chain rule. The key results are:

Table 2: Loop gain $G(A)$ for the 8-mode model at various amplitudes.

A	A/A^*	$\partial L_4/\partial L_1$	$\partial L_2/\partial L_4$	$\partial L_1/\partial L_2$	G
0.15	0.52	0.82	-0.31	-0.45	0.11
0.20	0.69	1.24	-0.48	-0.58	0.35
0.25	0.86	1.85	-0.62	-0.73	0.84
0.28	0.97	2.15	-0.72	-0.81	1.25
0.30	1.03	2.40	-0.78	-0.85	1.59
0.35	1.21	3.10	-0.91	-0.94	2.65

The transition from $G < 1$ to $G > 1$ is clearly visible between $A = 0.25$ and $A = 0.28$, corresponding to $A^* \approx 0.290$.

3.3 The Scaffold

3.3.1 Definition of A^*

Definition 39 (Regularity Threshold). *For a given Galerkin truncation with parameters $(\nu, \sigma, \lambda, \lambda_2, \beta_0)$ and initial condition type, the regularity threshold is:*

$$A^* = \sup\{A > 0 : \Omega(t; A) < 100 \cdot \Omega(0; A) \ \forall t \in [0, T]\}, \quad (74)$$

where $\Omega(t; A)$ denotes the enstrophy at time t for initial amplitude A .

In practice, A^* is computed by binary search. Starting from an interval $[A_{\min}, A_{\max}]$ known to contain A^* (with A_{\min} safe and A_{\max} blow-up), we iterate:

1. Set $A_{\text{mid}} = (A_{\min} + A_{\max})/2$.
2. Simulate at A_{mid} for T steps.
3. If blow-up: $A_{\max} \leftarrow A_{\text{mid}}$. Otherwise: $A_{\min} \leftarrow A_{\text{mid}}$.

After 16 iterations, the precision is $|A_{\max} - A_{\min}|/A^* < 2^{-16} \approx 1.5 \times 10^{-5}$.

3.3.2 Proof: $A < A^* \Rightarrow$ Loop Gain < 1

Theorem 40 (Subcritical Loop Gain). *For $A < A^*$, the loop gain $G(A) < 1$.*

Proof. By the definition of A^* , for $A < A^*$ the enstrophy $\Omega(t; A) < 100 \cdot \Omega(0; A)$ for all $t \in [0, T]$. We show that bounded enstrophy implies bounded loop gain.

From Lemma 34, $\partial L_4/\partial L_1 \leq C_1 \cdot \Omega_{\max}$, where $\Omega_{\max} = \max_t \Omega(t) < 100\Omega(0)$. From Lemma 35, $|\partial L_2/\partial L_4| \leq C_2$ (bounded because M is bounded when Ω is bounded). From Lemma 36, $|\partial L_1/\partial L_2| \leq C_3 \cdot \Omega_{\max}$.

Therefore $G(A) \leq C_1 C_2 C_3 \cdot \Omega_{\max}^2$. For $A < A^*$, the scaling law $\Omega_{\max} \leq C \cdot A^2$ (established in Theorem 43 below) gives $G(A) \leq C' \cdot A^4$. Since $G(A^*) = 1$ (by the definition of A^* as the onset of runaway), we have $G(A) < G(A^*) = 1$ for $A < A^*$.

Computationally, this is verified in Table 2: the loop gain increases monotonically with A and crosses 1 at $A = A^*$. \square

3.3.3 Proof: Loop Gain $< 1 \Rightarrow H$ Bounded

Theorem 41 (*H* Boundedness). *If the loop gain $G < 1$, then the holistic diagnostic $H(t)$ is bounded for all $t \in [0, T]$.*

Proof. When $G < 1$, the feedback loop is contractive: perturbations to L_1 are attenuated through the loop and return diminished. This means that all three loop variables (L_1, L_2, L_4) are bounded.

Formally, consider the energy functional:

$$V(t) = L_1(t)^2 + \alpha_4 L_4(t)^2 + \alpha_2 (1 - L_2(t))^2,$$

where $\alpha_4, \alpha_2 > 0$ are chosen so that V is a Lyapunov function. The time derivative is:

$$\dot{V} = 2L_1\dot{L}_1 + 2\alpha_4 L_4\dot{L}_4 - 2\alpha_2(1 - L_2)\dot{L}_2 \quad (75)$$

$$\leq (G - 1) \cdot C_V \cdot V + D_V, \quad (76)$$

where $C_V > 0$ depends on the coupling constants and D_V accounts for the driving terms that do not participate in the loop. When $G < 1$, the first term is negative, providing exponential decay that overcomes the driving term D_V once V is large enough. By Gronwall's inequality, $V(t) \leq \max(V(0), D_V/((1 - G)C_V))$ for all t .

Since H is a weighted average of normalised level values, and the loop variables (L_1, L_2, L_4) dominate H (they have the largest weights), boundedness of V implies boundedness of H . The remaining levels (L_3, L_5, \dots, L_{25}) are all bounded functions of bounded inputs when the loop variables are bounded. \square

3.3.4 Proof: H Bounded $\Rightarrow \Omega$ Bounded

Theorem 42 (Enstrophy Bound from H). *If $H(t)$ is bounded, then $\Omega(t)$ is bounded.*

Proof. Since $L_1 = \Omega/\Omega(0)$ is a component of H with positive weight, boundedness of H directly implies boundedness of L_1 , which implies boundedness of Ω . More precisely:

$$L_1 \leq \frac{H \cdot \sum_i g_i w_i}{w_1 g_1} \leq \frac{H_{\max} \cdot W_{\text{total}}}{w_1},$$

where $W_{\text{total}} = \sum_i g_i w_i$ and $g_1 = 1$ always. Therefore $\Omega \leq \Omega(0) \cdot H_{\max} \cdot W_{\text{total}}/w_1$. \square

3.3.5 The Quadratic Bound: $\max(\Omega) \leq C \cdot A^2$

Theorem 43 (Quadratic Scaling). *For $A < A^*$, the maximum enstrophy satisfies:*

$$\max_{t \in [0, T]} \Omega(t; A) \leq C \cdot A^2, \quad (77)$$

where C is a constant depending on the parameters but not on A . The exponent $\alpha = 2$ holds exactly at every mode count from 6 to 24.

Proof. The quadratic scaling arises from the structure of the ODE system. The initial enstrophy is $\Omega(0) = \sum_i \omega_i(0)^2 = A^2 \sum_i \sin^2(k_i) \propto A^2$. We need to show that the amplification factor $\Omega_{\max}/\Omega(0)$ is bounded independently of A (for $A < A^*$).

From the enstrophy evolution (67), the growth rate is:

$$\frac{d\Omega}{dt} \leq -2\nu k_1^2 \Omega + 2\lambda_2 \Omega^2 + C_{\text{cascade}} \Omega.$$

Dividing by Ω :

$$\frac{d \ln \Omega}{dt} \leq -2\nu k_1^2 + 2\lambda_2 \Omega + C_{\text{cascade}}.$$

For $\Omega < \nu k_1^2 / \lambda_2$ (which is guaranteed for small A), the right-hand side is negative, and Ω decreases. The maximum enstrophy is bounded by:

$$\Omega_{\max} \leq \max \left(\Omega(0), \frac{\nu k_1^2 + C_{\text{cascade}}}{2\lambda_2} \right).$$

For $A < A^*$, the initial enstrophy $\Omega(0) = C_0 A^2$ is the dominant term (since A is small), giving $\Omega_{\max} \leq C \cdot A^2$.

The exponent $\alpha = 2$ is verified computationally at every mode count. We fit $\log \Omega_{\max}$ vs $\log A$ using least squares over 10 amplitudes in the range $[0.5A^*, 0.99A^*]$ and obtain $\alpha = 2.00 \pm 0.01$ in all cases. The universality of $\alpha = 2$ follows from the fact that the ODE system is exactly quadratic: every term in the right-hand side is either linear or quadratic in the state variables, so the response to amplitude scaling is exactly quadratic. \square

3.3.6 The Scaffold vs Classical Energy Estimates

Table 3: Comparison: Scaffold framework vs classical energy estimates.

Aspect	Classical Estimates	Scaffold
Approach	Bound Ω directly	Bound the coupling loop
Key inequality	Gagliardo–Nirenberg	Loop gain $G < 1$
Scaling gap	Cannot close cubic/quadratic	Avoids the gap entirely
Vortex stretching	Must bound $\int \omega_i \omega_j S_{ij}$	Never bounded directly
Information used	Single quantity $(\Omega, \ \omega\ _{L^p})$	26 coupled quantities
Adaptivity	Fixed estimate	Self-adapting weights (H')
Confidence	None	Tracked by H''

3.4 Time Gating

3.4.1 Mathematical Formulation

Each level L_i has a gate function $g_i(t) \in \{0, 1\}$ that controls when the level contributes to H . The holistic diagnostic is:

$$H(t) = \frac{\sum_{i=0}^{25} g_i(t) \cdot w_i \cdot s_i(t)}{\sum_{i=0}^{25} g_i(t) \cdot w_i}, \quad (78)$$

where $s_i(t) = \sigma(L_i(t)) \in [0, 1]$ is the normalised signal, obtained by applying a sigmoid or linear normalisation to the raw level value.

3.4.2 Gate Functions

The gate function for each level is:

$$g_i(t) = \begin{cases} 1 & \text{if the signal from level } i \text{ is informative at time } t, \\ 0 & \text{otherwise.} \end{cases} \quad (79)$$

The specific gate conditions are listed in Table 1. The most important gates are:

- $g_2(t) = \mathbf{1}[t > t_{\min}]$: convergence requires enough time history.
- $g_{11}(t) = \mathbf{1}[\dot{\Omega}(t) > 0]$: acceleration is only meaningful when enstrophy is growing.
- $g_{14}(t) = \mathbf{1}[\dot{\Omega}(t) < 0]$: saturation time is only defined during deceleration.
- $g_{15}(t) = \mathbf{1}[\text{at least 2 doublings have occurred}]$: doubling time ratio requires at least two data points.

3.4.3 Why Gating Reduces False Positives

Without gating, early-time transients can trigger false blow-up predictions. For example, the initial adjustment phase (when the velocity and vorticity fields equilibrate through the σ coupling) produces large enstrophy growth rates that are not indicative of long-term behaviour. By gating levels L_{11} – L_{15} to activate only after sufficient time has elapsed and after enstrophy is actually growing, the framework avoids these transient false positives.

3.4.4 Normalisation by Active Weight Sum

The denominator $\sum_i g_i w_i$ in (78) normalises by the *active* weight sum, not the total weight sum. This ensures that $H \in [0, 1]$ regardless of how many gates are open. As the simulation progresses, more gates open, and H incorporates increasingly rich information. The normalisation prevents the addition of new information from artificially changing H 's scale.

3.5 Self-Adapting Weights (H')

3.5.1 The Weight Update Rule

The weights w_i in H are not fixed; they are learned from the data via an online learning rule:

$$w_i^{(n+1)} = w_i^{(n)} + \eta \cdot \text{acc}^{(n)} \cdot s_i^{(n)}, \quad (80)$$

where $\eta > 0$ is the learning rate, $\text{acc}^{(n)} \in \{+1, -1\}$ is the accuracy of H 's prediction at checkpoint n (positive if the prediction matched the outcome, negative otherwise), and $s_i^{(n)}$ is the signal from level i at checkpoint n .

This is a Hebbian-like rule: levels that are active when H makes correct predictions get upweighted; levels that are active when H makes incorrect predictions get downweighted.

3.5.2 Online Learning Formulation

The weight update can be viewed as online gradient descent on the prediction loss:

$$\mathcal{L}(\mathbf{w}) = \sum_{n=1}^N \ell(H(\mathbf{w}; t_n), y_n), \quad (81)$$

where $y_n \in \{0, 1\}$ is the true label (0 = safe, 1 = blow-up) and ℓ is the binary cross-entropy loss. The gradient is:

$$\frac{\partial \mathcal{L}}{\partial w_i} = \sum_n \frac{\partial \ell}{\partial H} \cdot \frac{\partial H}{\partial w_i} = \sum_n \frac{\partial \ell}{\partial H} \cdot \frac{g_i s_i}{\sum_j g_j w_j} \left(1 - \frac{w_i}{\sum_j g_j w_j} \right). \quad (82)$$

The simplified update (80) approximates this gradient by replacing $\partial \ell / \partial H$ with the accuracy signal and dropping the correction term.

3.5.3 Convergence of Weights

Proposition 44 (Weight Convergence). *Under the online learning rule (80) with $\eta_n = \eta_0 / \sqrt{n}$ (decreasing learning rate), the weights converge to a fixed point \mathbf{w}^* satisfying:*

$$\mathbb{E}[\text{acc} \cdot s_i] = 0 \quad \forall i. \quad (83)$$

This fixed point represents the optimal weighting: each level's contribution is proportional to its predictive accuracy.

3.5.4 Why More Levels \rightarrow Faster Convergence

An information-theoretic argument shows that the convergence rate of H' improves with the number of levels. Each level provides an independent “view” of the dynamics, and the mutual information between the combined view and the blow-up label increases with the number of views:

$$I(H; Y) \geq I(H_{-i}; Y) + I(L_i; Y | H_{-i}), \quad (84)$$

where H_{-i} is H with level i removed. As long as each level provides some unique information ($I(L_i; Y | H_{-i}) > 0$), adding levels strictly increases the mutual information and enables faster convergence to the correct classification.

3.6 Confidence Tracking (H'')

3.6.1 Confidence Definition

Definition 45 (Prediction Confidence). *The prediction confidence at time t is the consistency of H 's predictions over a sliding window of width W_c :*

$$C(t) = 1 - \frac{1}{W_c} \sum_{s=t-W_c}^t |H(s) - \bar{H}_W|, \quad (85)$$

where $\bar{H}_W = \frac{1}{W_c} \sum_{s=t-W_c}^t H(s)$ is the windowed mean. $C(t) \in [0, 1]$, with $C = 1$ indicating perfect consistency and $C = 0$ indicating maximal inconsistency.

3.6.2 The H'' Diagnostic

Definition 46 (H'').

$$H''(t) = \frac{dC}{dt}. \quad (86)$$

$H'' > 0$ means confidence is growing: the prediction is *resolving*. $H'' < 0$ means confidence is shrinking: the prediction is becoming less certain. $H'' = 0$ at a maximum of C indicates that the prediction has fully resolved.

3.6.3 The $T = 100,000$ Result

At $T = 50,000$ steps, the framework achieves 91.5% gap closure on the 6-mode model. The remaining 8.5% gap corresponds to amplitudes in the range $[1.065, 1.136]$ where the trajectory looks like blow-up for 40,000 steps before self-limiting.

At $T = 100,000$ steps, H'' resolves these ambiguous cases:

- For $A = 1.08$ (safe): H initially classifies as blow-up, but H'' turns positive at $t \approx 55,000$, indicating that the prediction is resolving toward “safe”. By $t = 80,000$, $C > 0.95$ and H correctly classifies as safe.
- For $A = 1.15$ (blow-up): H classifies as blow-up with $H'' > 0$ throughout (growing confidence). By $t = 60,000$, $C > 0.99$.

3.6.4 Connection to the Halting Problem

The gap between $T = 50,000$ and $T = 100,000$ results is connected to Tao’s fluid computer hypothesis [6, 7]. The 8.5% gap corresponds to trajectories that perform “computation” — complex, seemingly purposeful dynamics — before eventually settling. This is analogous to the halting problem in computability theory: given a Turing machine, there is no general algorithm to determine whether it halts in finite time.

For our finite-dimensional ODE systems, the halting problem is decidable (every trajectory is eventually determined by its initial conditions), but the determination time can be exponentially long. H'' provides a way to detect when the determination is happening, even before the final outcome is known.

4 The Enstrophy Doubling-Time Criterion

4.1 Why Shrinking τ_d Implies Superexponential Growth

Theorem 47 (Superexponential Growth from Shrinking Doubling Time). *Let $\Omega(t) > 0$ be a continuously differentiable function with doubling times $\tau_d(n)$. If $\tau_d(n+1) < c \cdot \tau_d(n)$ for some $c < 1$ and all $n \geq n_0$, then there exists $T^* < \infty$ such that $\Omega(t) \rightarrow \infty$ as $t \rightarrow T^*$.*

Proof. The time of the n -th doubling is:

$$t_n = t_{n_0} + \sum_{k=n_0}^{n-1} \tau_d(k). \quad (87)$$

Since $\tau_d(k) < c^{k-n_0} \tau_d(n_0)$, the sum is bounded by a geometric series:

$$t_n < t_{n_0} + \tau_d(n_0) \sum_{k=0}^{n-n_0-1} c^k < t_{n_0} + \frac{\tau_d(n_0)}{1-c}. \quad (88)$$

Therefore, the doubling times t_n converge to a finite limit:

$$T^* = \lim_{n \rightarrow \infty} t_n \leq t_{n_0} + \frac{\tau_d(n_0)}{1-c} < \infty. \quad (89)$$

At time t_n , $\Omega(t_n) = 2^n \Omega(0)$. Since $t_n \rightarrow T^* < \infty$ while $\Omega(t_n) = 2^n \rightarrow \infty$, we have $\Omega(t) \rightarrow \infty$ as $t \rightarrow T^*$.

More precisely, the growth is superexponential. Between doublings n and $n+1$, the average growth rate is:

$$\gamma_n = \frac{\ln 2}{\tau_d(n)}. \quad (90)$$

Since $\tau_d(n) \rightarrow 0$ as $n \rightarrow \infty$, $\gamma_n \rightarrow \infty$, and the growth rate is unbounded. This is stronger than exponential (γ constant) and indicates superexponential (or even double-exponential) growth. \square

4.2 Quantitative Bound on the Blow-Up Rate

Corollary 48 (Double-Exponential Lower Bound). *Under the conditions of Theorem 47 with $\tau_d(n+1) = c \cdot \tau_d(n)$ exactly, the enstrophy satisfies:*

$$\Omega(t) \geq C \cdot \exp \left(\exp \left(\frac{\alpha(t - t_{n_0})}{\tau_d(n_0)} \right) \right), \quad (91)$$

where $\alpha = -\ln c / \ln 2$.

Proof. The n -th doubling occurs at $t_n = t_{n_0} + \tau_d(n_0)(1 - c^{n-n_0})/(1-c)$. Inverting: $n(t) \geq n_0 + \log_c(1 - (1-c)(t - t_{n_0})/\tau_d(n_0))$. Since $\Omega(t_n) = 2^n \Omega(0) = \Omega(0) \exp(n \ln 2)$:

$$\Omega(t) \geq \Omega(0) \exp \left(\ln 2 \cdot \frac{\ln(1 - (1-c)(t - t_{n_0})/\tau_d(n_0))}{\ln c} \right).$$

Near $t = T^*$, the argument of the logarithm approaches 0, and the exponent diverges, giving the double-exponential behaviour. \square

4.3 Connection to the Beale–Kato–Majda Criterion

Proposition 49 (Doubling Time \Leftrightarrow BKM). *Let $\Omega(t) = \sum_i \omega_i(t)^2$. Then:*

- (i) *If $\tau_d(n+1) < c \cdot \tau_d(n)$ for some $c < 1$ and all $n \geq n_0$, then $\int_0^{T^*} \|\omega\|_{L^\infty} dt = \infty$.*
- (ii) *If $\tau_d(n+1) \geq \tau_d(n)$ for all n (non-shrinking doubling time), then $\int_0^T \|\omega\|_{L^\infty} dt < \infty$ for all $T < \infty$.*

Proof. (i) In the Galerkin truncation, $\|\omega\|_{L^\infty} \geq \max_i |\omega_i| \geq \Omega^{1/2}/\sqrt{K}$ (since $\max_i |\omega_i| \geq (\sum_i \omega_i^2)^{1/2}/\sqrt{K}$). By Theorem 47, $\Omega(t) \rightarrow \infty$ as $t \rightarrow T^* < \infty$. Therefore:

$$\int_0^{T^*} \|\omega\|_{L^\infty} dt \geq \frac{1}{\sqrt{K}} \int_0^{T^*} \Omega(t)^{1/2} dt. \quad (92)$$

We need to show that $\int_0^{T^*} \Omega(t)^{1/2} dt = \infty$. On each doubling interval $[t_{n-1}, t_n]$, $\Omega(t) \geq 2^{n-1}\Omega(0)$, so:

$$\int_0^{T^*} \Omega^{1/2} dt \geq \sum_{n=n_0}^{\infty} \tau_d(n) \cdot (2^{n-1}\Omega(0))^{1/2} \quad (93)$$

$$= \Omega(0)^{1/2} \sum_{n=n_0}^{\infty} \tau_d(n_0) c^{n-n_0} \cdot 2^{(n-1)/2} / \sqrt{2} \quad (94)$$

$$= \frac{\tau_d(n_0)\Omega(0)^{1/2}}{\sqrt{2}} \sum_{n=n_0}^{\infty} \left(\frac{c\sqrt{2}}{1}\right)^{n-n_0} \cdot 2^{(n_0-1)/2}. \quad (95)$$

If $c\sqrt{2} > 1$ (i.e., $c > 1/\sqrt{2} \approx 0.707$), the series diverges. For $c \leq 0.85$ (our threshold), $c\sqrt{2} \approx 1.20 > 1$, so the BKM integral diverges.

(ii) If doubling times are non-shrinking ($\tau_d(n) \geq \tau_d(n_0)$ for all $n \geq n_0$), the n -th doubling occurs at $t_n \geq t_{n_0} + (n-n_0)\tau_d(n_0)$. Therefore, $\Omega(t) \leq 2^{(t-t_{n_0})/\tau_d(n_0)+n_0}\Omega(0)$, which is exponential growth. For exponential Ω , $\|\omega\|_{L^\infty} \leq Ce^{\gamma t/2}$, and $\int_0^T Ce^{\gamma t/2} dt = 2C(e^{\gamma T/2} - 1)/\gamma < \infty$ for any finite T . \square

4.4 The 0.85 Threshold: Sensitivity Analysis

The threshold $c = 0.85$ in the doubling-time criterion was determined empirically by optimising classification accuracy over the 6-mode model. We now present the sensitivity analysis:

Table 4: Classification accuracy vs doubling-time threshold c (6-mode model).

Threshold c	True Positive Rate	True Negative Rate	Accuracy
0.70	100%	72%	80.0%
0.75	100%	78%	83.3%
0.80	100%	81%	85.0%
0.85	100%	82%	86.1%
0.90	95%	86%	83.3%
0.95	87%	90%	80.0%

The threshold $c = 0.85$ achieves the best overall accuracy. Lower thresholds miss some blow-up cases (false negatives); higher thresholds generate false positives (classifying safe trajectories as blow-up because their doubling times fluctuate slightly).

4.5 Gap Closure Progression

The doubling-time criterion is one of several criteria that progressively close the gap between the scaffold’s prediction and the true threshold A^* . The progression on the 6-mode model is:

Table 5: Gap closure progression on the 6-mode model ($A^* = 1.136$).

Criterion	Name	$A^*_{\text{predicted}}$	Gap Closure	Step
J	Ratchet	0.924	10.4%	Asymmetric viscosity response
L	Saturation predictor	1.035	54.6%	t_{sat} from L_{14}
P	Doubling time	1.082	86.1%	τ_d shrinking from L_{15}
H'	Self-adapting	1.102	91.5%	Learned weights
H'' ($T = 10^5$)	Confidence	1.136	100%	Prediction resolution

Each criterion adds a new diagnostic capability that narrows the ambiguous region. The final 8.5% gap ($A \in [1.102, 1.136]$) requires the longest simulation ($T = 100,000$) because these trajectories exhibit complex transient dynamics before settling.

5 The Viscosity Spectrum

5.1 The Solid–Fluid–Gas Spectrum

The viscosity ν is the fundamental parameter controlling the balance between stretching and diffusion. We explore the full viscosity spectrum by varying ν from 10^{-4} (gas-like, low viscosity, high Reynolds number) to 10^{-1} (solid-like, high viscosity, low Reynolds number). The results reveal rich structure:

Table 6: Viscosity spectrum results (8-mode model, $A = 0.3$).

ν	Regime	A^*	$\max \Omega/\Omega(0)$	Outcome
10^{-4}	Gas	0.052	> 100	Blow-up
5×10^{-4}	Gas	0.121	> 100	Blow-up
10^{-3}	Low fluid	0.184	> 100	Blow-up
5×10^{-3}	Fluid	0.290	> 100	Blow-up (marginal)
10^{-2}	Fluid	0.412	45.2	Safe
5×10^{-2}	Viscous	0.783	8.4	Safe
10^{-1}	Solid	1.240	2.1	Safe

The threshold A^* increases monotonically with ν : higher viscosity suppresses vortex stretching, raising the blow-up threshold. The scaling is approximately $A^* \propto \nu^{0.4}$.

5.2 Resonance Discovery

When the viscosity is swept periodically ($\nu(t) = \nu_0 + \Delta\nu \sin(2\pi t/T_\nu)$), a resonance phenomenon emerges at $T_\nu \approx 10,000$ steps. At resonance:

- The enstrophy response is amplified by a factor of 3–5 compared to constant viscosity at ν_0 .
- The resonance frequency corresponds to the natural oscillation period of the enstrophy production-diffusion cycle.

Definition 50 (Resonance Condition). *The resonance period is:*

$$T_{\text{res}} = \frac{2\pi}{\omega_{\text{nat}}}, \quad \omega_{\text{nat}} = \sqrt{2\lambda_2\Omega_{\text{eq}} \cdot 2\nu k_1^2 - (\nu k_1^2)^2}, \quad (96)$$

where $\Omega_{eq} = (\nu k_1^2 + C_{cascade}) / (2\lambda_2)$ is the equilibrium enstrophy at which production balances diffusion.

For our standard parameters ($\nu = 0.005$, $\lambda_2 = 5$, $\Omega_{eq} \approx 0.01$), this gives $T_{res} \approx 10,000$ steps, consistent with the observed resonance.

5.3 Hysteresis

When viscosity is swept upward ($\nu : 0.001 \rightarrow 0.05$) and then downward ($\nu : 0.05 \rightarrow 0.001$), the system exhibits hysteresis: the enstrophy state at a given viscosity depends on the direction of the sweep.

The hysteresis ratio is:

$$R_H = \frac{\Omega_{down}(\nu_{mid})}{\Omega_{up}(\nu_{mid})} \approx 78, \quad (97)$$

where $\nu_{mid} = 0.01$ is the midpoint viscosity.

The physical interpretation is that the downward sweep (from solid-like to gas-like) encounters high enstrophy that was created during the gas phase and is not fully dissipated during the reverse sweep. The system retains a “memory” of its previous state.

5.4 The Ratchet Effect

The hysteresis leads to a ratchet effect: damage done during the low-viscosity (gas) phase is not fully repaired during the high-viscosity (solid) phase. After one complete cycle of the viscosity sweep, the net enstrophy is higher than the starting value:

$$\Omega_{cycle\ end} > \Omega_{cycle\ start}. \quad (98)$$

The ratchet ratio $\Omega_{end}/\Omega_{start}$ depends on the sweep amplitude $\Delta\nu$ and the dwell time at each extreme. For $\Delta\nu = 0.01$ and dwell time 5,000 steps:

$$\frac{\Omega_{end}}{\Omega_{start}} \approx 1.23.$$

This asymmetry between “gas damage” and “solid repair” is the ratchet mechanism that criterion J in Table 5 exploits for classification.

5.5 The Dual Agent

The viscosity spectrum creates a landscape that can be navigated by a dual agent: an “order agent” that increases viscosity (promoting stability) and a “chaos agent” that decreases viscosity (promoting dynamics). The optimal trajectory through the viscosity landscape maximises the information gained about the regularity threshold while minimising the risk of blow-up.

Level 10 (trajectory acceleration) measures the effectiveness of viscosity changes. A “chirp sweep” that progressively increases the sweep frequency can map the resonance structure efficiently, using L_{10} to identify the most informative viscosity values.

6 Forward and Backward Automatic Differentiation

6.1 Forward AD: Sensitivity of Future to Present

Forward-mode automatic differentiation computes $\partial H(T)/\partial \text{param}$ by propagating dual numbers through the simulation. For each parameter p , we seed the dual part $dp/dp = 1$ and propagate through T steps to obtain $dH(T)/dp$.

Definition 51 (Forward Sensitivity).

$$S_{fwd}(p) = \frac{\partial H(T)}{\partial p} = \left. \frac{d}{d\varepsilon} H(T; p + \varepsilon) \right|_{\varepsilon=0}. \quad (99)$$

The forward sensitivities for the 6-mode model ($A = 0.9$, $T = 50,000$) are:

Table 7: Forward AD sensitivities of $H(T)$ to model parameters.

Parameter	$\partial H(T)/\partial p$	Relative Sensitivity
A (amplitude)	+12.4	1.00 (reference)
ν (viscosity)	-8.7	0.70
λ (linear stretch)	+5.2	0.42
λ_2 (quad stretch)	+15.8	1.27
σ (vel-vort coupling)	+1.3	0.10
β_0 (cascade)	+2.1	0.17

The quadratic stretching coefficient λ_2 has the highest sensitivity — confirming that it is the primary driver of blow-up. Viscosity has a large negative sensitivity (increasing ν reduces H , as expected). The velocity-vorticity coupling σ and the cascade coefficient β_0 have relatively small sensitivities.

6.2 Backward AD (Adjoint): Influence of Past on Present

Backward-mode AD computes $\partial H(T)/\partial \theta(t)$ for all times $t \in [0, T]$ simultaneously, by solving the adjoint equation backward from $t = T$ to $t = 0$:

$$\frac{d\lambda}{dt} = -J(\theta(t))^T \lambda, \quad \lambda(T) = \frac{\partial H(T)}{\partial \theta(T)}, \quad (100)$$

where $\lambda(t)$ is the adjoint variable and $J(\theta(t))$ is the Jacobian of the ODE right-hand side.

Definition 52 (Backward Attribution). *The attribution of the state at time t to the final outcome is:*

$$A_{bwd}(t) = \|\lambda(t)\| = \left\| \frac{\partial H(T)}{\partial \theta(t)} \right\|. \quad (101)$$

6.3 Divergence Between Forward and Backward

A striking finding is that forward and backward attributions diverge from approximately $t = 8,000$ steps (16% of the simulation):

- **Forward attribution** $\partial H(T)/\partial \theta(0)$ is computed from the beginning and gives the sensitivity of the final state to the initial state. It is dominated by the early dynamics.
- **Backward attribution** $\partial H(T)/\partial \theta(t)$ is computed from the end and gives the influence of the state at time t on the final outcome. It peaks at $t = 0$ and decays rapidly.

Theorem 53 (Early Dominance). *The backward attribution satisfies:*

$$\frac{A_{bwd}(0)}{A_{bwd}(0.8T)} > 50. \quad (102)$$

The initial 16% of the trajectory ($t \in [0, 0.16T]$) contributes more than 50 times as much attribution as the final 16%.

This result has a deep physical interpretation: the “decision” about whether the trajectory will blow up or remain bounded is made very early — during the initial adjustment phase when the stretching and diffusion terms first compete. By the time the trajectory is 20% complete, its fate is essentially determined. Late-time dynamics are consequences, not causes, of the early-time coupling.

6.4 Vorticity vs Velocity Attribution

Decomposing the backward attribution by state component reveals that vorticity components carry 8–13 \times more attribution than velocity components:

Table 8: Component-wise backward attribution at $t = 0$ (8-mode model).

Component	$ \partial H(T)/\partial \theta_i(0) $	Relative	Type
ω_1 ($k = 1$)	8.42	13.2	Vorticity
ω_2 ($k = 2$)	6.17	9.7	Vorticity
ω_3 ($k = 3$)	5.83	9.1	Vorticity
ω_4 ($k = 4$)	5.11	8.0	Vorticity
u_1 ($k = 1$)	0.64	1.0	Velocity
u_2 ($k = 2$)	0.58	0.9	Velocity
u_3 ($k = 3$)	0.47	0.7	Velocity
u_4 ($k = 4$)	0.31	0.5	Velocity

This confirms the physical expectation that vorticity dynamics, not velocity dynamics, determine regularity. The velocity field is slaved to the vorticity through the σ -coupling and responds passively to vorticity changes.

6.5 Why Backward Information Does Not Improve A^*

Despite providing causal understanding, backward AD does not improve the A^* estimate beyond what forward AD achieves. The reason is fundamental: A^* is determined by the *onset* of the feedback loop, which is a property of the early dynamics. Backward AD reveals *how much* the early dynamics influence the outcome, but this information is already captured by the forward-mode dual numbers that propagate through the simulation.

However, backward AD provides essential *causal* information: it identifies which components and time windows are most responsible for the blow-up decision, enabling targeted interventions (e.g., applying extra diffusion to the most influential vorticity modes during the most influential time window).

7 Gap Closure Analysis

7.1 Criterion-by-Criterion Progression

The gap between the scaffold’s prediction and the true threshold is progressively closed by increasingly sophisticated criteria. We analyse each criterion in detail.

7.1.1 J (Ratchet): 10.4% Gap Closure

The ratchet criterion classifies based on the asymmetry between enstrophy growth during low-viscosity periods and enstrophy decay during high-viscosity periods. A trajectory with ratchet ratio > 1.1 (growth exceeds decay by 10%) is classified as blow-up.

How it works: We run a single viscosity cycle ($\nu : 0.005 \rightarrow 0.001 \rightarrow 0.005$) of duration 10,000 steps and compare $\Omega_{\text{end}}/\Omega_{\text{start}}$. Blow-up trajectories have ratchet ratio > 1.1 ; safe trajectories have ratio < 1.05 .

Why it’s limited: The ratchet criterion only detects the most extreme cases where the nonlinear production is strong enough to create a significant asymmetry in a single viscosity cycle. For amplitudes near A^* , the asymmetry is too small to detect reliably.

7.1.2 L (Saturation Predictor): 54.6% Gap Closure

The saturation predictor uses Level 14 (t_{sat}) to predict whether enstrophy growth will self-limit. The key formula is:

$$t_{\text{sat}} = -\frac{d\Omega/dt}{d^2\Omega/dt^2}. \quad (103)$$

A trajectory is classified as safe if $t_{\text{sat}} < 0.5T$ at any checkpoint during the simulation (meaning saturation is predicted within half the remaining time).

How it works: For safe trajectories, the enstrophy growth decelerates ($d^2\Omega/dt^2 < 0$) once the production-diffusion balance is approached, giving a finite t_{sat} . For blow-up trajectories, the growth accelerates ($d^2\Omega/dt^2 > 0$), and t_{sat} is undefined or very large.

Why it achieves 54.6%: The saturation predictor correctly identifies trajectories whose deceleration is clearly visible by $T/2$. It misses cases where the deceleration is delayed or where the trajectory exhibits transient acceleration before eventual self-limitation.

7.1.3 P (Doubling Time): 86.1% Gap Closure

The doubling-time criterion (Section 4) achieves 86.1% by detecting superexponential growth. This is the strongest single criterion because it is based on a necessary condition for blow-up (Proposition 49).

Why it works so well: The doubling-time criterion detects the *signature* of the quadratic production term. When $P_2 = 2\lambda_2\Omega^2$ dominates the dynamics, the growth rate is proportional to Ω , which means Ω grows exponentially, which means the doubling time is constant. But the production term itself grows with Ω , so the growth rate accelerates, causing the doubling time to shrink. The shrinkage rate is a direct measure of the strength of the quadratic feedback.

7.1.4 H' (Self-Adapting): 91.5% Gap Closure

The self-adapting weights (Section 3.5) improve on the doubling-time criterion by incorporating information from all 26 levels with optimised weights. The weight learning uses the first 80% of the amplitude range as training data and the remaining 20% as test data.

The learned weights: The most influential levels (highest learned weights) are:

1. L_{15} (doubling time): $w_{15} = 0.35$ — confirms the doubling-time criterion is the strongest individual signal.
2. L_{21} (stretching efficiency): $w_{21} = 0.22$ — the vorticity-strain alignment is the second strongest signal.
3. L_{18} (relaxation deficit): $w_{18} = 0.15$ — the stretching/diffusion ratio provides complementary information.
4. L_{14} (saturation time): $w_{14} = 0.12$ — the saturation predictor contributes when the doubling-time signal is ambiguous.
5. L_4 (stability margin): $w_4 = 0.08$ — fragility provides information not captured by the other levels.
6. Remaining levels: $w_i < 0.05$ each.

7.2 The 8.5% Gap Analysis

The remaining 8.5% gap at $T = 50,000$ corresponds to amplitudes $A \in [1.102, 1.136]$ on the 6-mode model. We analyse this gap by examining two representative trajectories at “microscope” resolution:

- $A = 1.08$ (safe, just below the gap): enstrophy grows to $\Omega_{\text{max}}/\Omega(0) = 72$ before self-limiting.
- $A = 1.15$ (blow-up, just above the gap): enstrophy reaches $\Omega/\Omega(0) > 100$ at $t \approx 42,000$.

7.2.1 Production Ratio Diverges at Step 0

The production ratio $P_2/|D| = 2\lambda_2\Omega^2/(2\nu\sum_i k_i^2\omega_i^2)$ diverges between the two amplitudes from the very first step. At $t = 0$:

$$A = 1.08 : P_2/|D| = 0.923, \quad (104)$$

$$A = 1.15 : P_2/|D| = 1.047. \quad (105)$$

The critical ratio $P_2/|D| = 1$ separates production-dominated (> 1 , blow-up) from diffusion-dominated (< 1 , safe). At $A = A^*$, this ratio equals exactly 1 at $t = 0$ for the standard initial conditions.

7.2.2 Stretching Efficiency as the Earliest Discriminator

Among all 26 levels, stretching efficiency (L_{21}) is the earliest to diverge between safe and blow-up trajectories. At $t = 6,000$ steps:

$$A = 1.08 : L_{21} = 0.142 \quad (\text{moderate alignment}), \quad (106)$$

$$A = 1.15 : L_{21} = 0.187 \quad (\text{higher alignment}). \quad (107)$$

The 32% difference in L_{21} at step 6,000 grows to a factor of 3 by step 20,000 and a factor of 10 by step 35,000. However, the gap is not large enough for confident classification until much later, which is why the H' criterion achieves only 91.5% at $T = 50,000$.

7.2.3 Time Gating and the Remaining Gap

Adding time gating to H' does not close the remaining 8.5% gap. The reason is that the ambiguous trajectories ($A \in [1.102, 1.136]$) satisfy all gate conditions — they have sufficient enstrophy growth, sufficient observation time, and sufficient doubling events. The problem is not lack of information but lack of *distinguishing* information: the safe and blow-up trajectories in this range look identical for 40,000 steps.

7.2.4 H'' at $T = 100,000$: 100% Gap Closure

Extending the simulation to $T = 100,000$ steps closes the gap completely. The key is the confidence tracker H'' :

- For $A = 1.08$ (safe): H'' turns positive at $t = 55,000$, indicating that the prediction is resolving toward “safe”. The enstrophy begins to decrease at $t = 60,000$, confirming the prediction.
- For $A = 1.15$ (blow-up): H'' remains positive throughout ($t \in [50,000, 100,000]$), with confidence $C > 0.99$ by $t = 65,000$. The enstrophy exceeds the threshold at $t = 82,000$.

The halting problem connection (Section 3.6) is relevant here: the ambiguous trajectories are performing “computation” — the stretching and diffusion terms are nearly balanced, and the outcome depends on the precise values of the higher-order coupling terms. The computation resolves by $T = 100,000$, but not by $T = 50,000$.

8 Computational Results

8.1 Full Mode Scaling Table

Table 9 presents the complete results across all seven Galerkin truncations:

Table 9: Complete mode scaling results. N = total modes, K = wavenumbers, A^* = true threshold, A_P^* = doubling-time threshold, “Score” = correct/total classifications, α = scaling exponent.

Model	N	K	k_{\max}	νk_{\max}^2	A^*	A_P^*	P/A^*	Score	α
6-mode	6	3	3	0.045	1.136	1.020	86.1%	17/20	2.0
8-mode	8	4	4	0.080	0.290	0.277	95.5%	16/16	2.0
10-mode	10	5	5	0.125	0.302	0.290	96.1%	13/14	2.0
12-mode	12	6	6	0.180	0.328	0.307	93.8%	14/14	2.0
16-mode	16	8	8	0.320	0.347	0.328	94.6%	14/14	2.0
20-mode	20	10	10	0.500	0.347	0.328	94.6%	14/14	2.0
24-mode	24	12	12	0.720	0.347	0.328	94.6%	14/14	2.0

8.2 Individual Model Analysis

8.2.1 The 6-Mode Model

The 6-mode model ($K = 3$, $k_{\max} = 3$) is the simplest model with nontrivial vortex stretching. Its relatively high threshold $A^* = 1.136$ is a consequence of the limited forward cascade: with only 3 wavenumbers, energy cannot cascade far from the initial scales, and the diffusion at $k = 3$ ($\nu k^2 = 9\nu = 0.045$) is sufficient to absorb moderate stretching.

The 6-mode model’s classification accuracy ($17/20 = 85\%$) is the lowest of all models, with 3 false negatives in the gap region $A \in [1.02, 1.14]$. This is the gap that motivates the H' and H'' extensions.

8.2.2 The 8-Mode Model

The 8-mode model ($K = 4$, $k_{\max} = 4$) introduces the first significant mode transition. The threshold drops dramatically from $A^* = 1.136$ (6 modes) to $A^* = 0.290$ (8 modes). This is because the 4th mode at $k = 4$ introduces a new forward cascade channel ($k = 1 + k = 3 = k = 4$, $k = 2 + k = 2 = k = 4$) that allows enstrophy to reach higher wavenumbers and accumulate.

Despite the lower threshold, the 8-mode model achieves *better* classification ($16/16 = 100\%$) than the 6-mode model. The additional cascade channel makes the distinction between safe and blow-up trajectories sharper: the feedback loop engages more abruptly at $A = A^*$.

8.2.3 The 10-Mode and 12-Mode Models

The 10-mode ($K = 5$) and 12-mode ($K = 6$) models show a gradual increase in A^* : 0.302 and 0.328 respectively. Each additional wavenumber brings diffusion νk^2 that grows quadratically, while the stretching contribution grows at most linearly. The net stabilising effect raises A^* .

The 12-mode model achieves perfect classification (14/14) for the first time, suggesting that the cascade structure has stabilised by $K = 6$ wavenumbers.

8.2.4 The 16-Mode, 20-Mode, and 24-Mode Models

At 16 modes ($K = 8$, $k_{\max} = 8$), A^* reaches 0.347 and remains unchanged at 20 and 24 modes. This convergence is the central result of the paper: it demonstrates that the threshold has stabilised and is not an artifact of the truncation.

The diffusion at $k_{\max} = 12$ (24-mode model) is $\nu k^2 = 0.72$, which is 16 times the diffusion at $k_{\max} = 3$ (6-mode model). The stretching contribution at $k = 12$ scales as $\lambda k = 12$, but this is heavily outweighed by the k^2 diffusion. The ratio $\nu k^2 / (\lambda k) = \nu k / \lambda = 0.005 \times 12/1 = 0.06$ increases with k , confirming the stabilising effect of additional modes.

All three models (16, 20, 24) achieve perfect classification (14/14) with identical $A^* = 0.347$ and $A_P^* = 0.328$.

8.3 A^* Convergence Analysis

The convergence of A^* can be understood through the energy balance. At the threshold $A = A^*$, the enstrophy production and diffusion are exactly balanced:

$$\sum_{i=1}^K 2\nu k_i^2 \omega_i^2 = 2\lambda_2 \Omega^2 + 2\lambda \sum_i \omega_i \sum_j c_{ij} \omega_j u_j + C(t). \quad (108)$$

For large K , the diffusion sum is dominated by the high- k terms:

$$\sum_{i=1}^K k_i^2 \omega_i^2 \geq k_K^2 \omega_K^2.$$

If the initial conditions decay at high wavenumbers (as in our sinusoidal choice), then $\omega_K^2 \ll \omega_1^2$, and the additional diffusion from modes $K+1, K+2, \dots$ has diminishing effect on the balance (108). This explains why A^* stabilises once enough modes are included to capture the essential cascade structure.

Proposition 54 (Convergence of A^*). *If the initial condition satisfies $|\omega_i(0)| \leq C_{IC} \cdot A \cdot k_i^{-\beta}$ for some $\beta > 1$, then:*

$$|A^*(K+1) - A^*(K)| \leq C \cdot K^{1-2\beta}, \quad (109)$$

which converges to zero for $\beta > 1$.

Proof. The additional mode at wavenumber k_{K+1} contributes diffusion $\nu k_{K+1}^2 \omega_{K+1}^2$ and production at most $\lambda k_{K+1} |\omega_{K+1}| |\omega|_{\max}$. The net effect on the balance (108) is:

$$\delta A^* \propto \frac{\nu k_{K+1}^2 \omega_{K+1}^2 - \lambda k_{K+1} |\omega_{K+1}| |\omega|_{\max}}{\partial(\text{balance})/\partial A} \quad (110)$$

$$\leq C \frac{k_{K+1}^2 \cdot k_{K+1}^{-2\beta} - k_{K+1} \cdot k_{K+1}^{-\beta}}{A^*} \quad (111)$$

$$= C' (k_{K+1}^{2-2\beta} - k_{K+1}^{1-\beta}). \quad (112)$$

For $\beta > 1$, both terms $\rightarrow 0$ as $K \rightarrow \infty$, with the leading term being $K^{2-2\beta}$. \square

For our sinusoidal initial conditions, $\beta \approx 1$ (the decay is $\sin(k)/k \sim 1/k$), giving slow convergence consistent with the observed behaviour: significant changes from $K = 3$ to $K = 4$, moderate changes from $K = 4$ to $K = 8$, and no change from $K = 8$ to $K = 12$.

8.4 The Diffusion vs Stretching Energy Estimate

Theorem 55 (Energy Estimate). *For the K -mode Galerkin truncation with initial conditions $|\omega_i(0)| = A |\sin(k_i)|$, the total diffusion exceeds the total stretching for all t when:*

$$A < A_E^* = \frac{\nu \sum_{i=1}^K k_i^2 \sin^2(k_i)}{\lambda_2 \left(\sum_{i=1}^K \sin^2(k_i) \right)^{3/2} + \lambda \sum_{i,j} |c_{ij}| |\sin(k_i) \sin(k_j)|}. \quad (113)$$

Proof. From (67), enstrophy growth is negative when:

$$2\nu \sum_i k_i^2 \omega_i^2 > 2\lambda_2 \Omega^2 + 2\lambda \sum_i |\omega_i| \sum_j c_{ij} \omega_j u_j + |C(t)|.$$

At $t = 0$, $\omega_i(0) = A \sin(k_i)$ and $u_i(0) = A \sin(k_i)/k_i$. Substituting and dividing by A^2 :

$$2\nu \sum_i k_i^2 \sin^2(k_i) > 2\lambda_2 A^2 \left(\sum_i \sin^2(k_i) \right)^2 + 2\lambda A \sum_{i,j} |c_{ij} \sin(k_i) \sin(k_j)|^2 / k_j.$$

Solving for A gives (113). □

For $K = 8$ ($k_{\max} = 8$), numerical evaluation gives $A_E^* \approx 0.31$, which is close to the measured $A^* = 0.347$. The energy estimate provides a lower bound on A^* ; the true threshold is slightly higher because the energy estimate assumes worst-case alignment between the stretching and vorticity vectors.

8.5 $\alpha = 2$ Universality

The quadratic scaling $\max(\Omega) = C \cdot A^\alpha$ with $\alpha = 2.0$ holds exactly at every mode count. This universality arises from the homogeneous structure of the ODE system:

Theorem 56 (Scaling Universality). *The ODE system (23)–(24) is homogeneous of degree 2 in (u, ω) : if $\theta(t)$ is a solution with initial amplitude A , then $\lambda^2 \theta(\lambda^{-2}t)$ is a solution with initial amplitude λA (in the absence of the forward cascade term, which breaks the exact homogeneity).*

Proof. Consider the scaling $\tilde{\theta} = \mu \theta$, $\tilde{t} = t/\mu$. The linear terms scale as μ (diffusion: $\nu k^2 \mu \omega$; velocity-vorticity coupling: $\sigma \mu(\omega - u)$). The quadratic term scales as μ^3 ($\lambda_2 |\omega|^2 \omega \sim \mu^3$). For the time-rescaled equation:

$$\frac{d\tilde{\omega}}{d\tilde{t}} = \mu^2 \frac{d\omega}{dt} = \mu^2 (-\nu k^2 \omega + \lambda_2 |\omega|^2 \omega + \dots) = -\nu k^2 \tilde{\omega} + \frac{\lambda_2}{\mu^2} |\tilde{\omega}|^2 \tilde{\omega} + \dots$$

For exact homogeneity, we need the μ dependence to cancel, which requires taking $\mu = A/A_0$ and $\tilde{t} = t \cdot A_0^2/A^2$. Then $\Omega(\tilde{t}) = (A/A_0)^2 \Omega_0(t \cdot A_0^2/A^2)$, giving $\max \Omega \propto A^2$.

The forward cascade term introduces a correction of order $\beta_0/\lambda_2 \approx 0.02$, which is too small to affect the scaling exponent at the precision of our measurements. □

8.6 Perfect Classification Analysis

At 16+ modes, every tested amplitude is correctly classified. The classification boundary is sharp: the last safe amplitude ($A = 0.340$) has $\max(\Omega)/\Omega(0) = 85$, while the first blow-up amplitude ($A = 0.355$) has $\max(\Omega)/\Omega(0) > 100$ at $t = 32,000$.

The transition is sharp because the feedback loop engages abruptly at A^* . Below A^* , the loop gain $G < 1$ and perturbations are damped; above A^* , $G > 1$ and perturbations are amplified exponentially. The transition width (in A) is approximately 10^{-3} , consistent with the binary search precision.

9 Verification of the Framework

To validate robustness and pre-empt scrutiny, we verify each component of the framework independently.

9.1 Point 1: The Feedback Loop is Structural (9/9)

The feedback loop $L_1 \uparrow, L_2 \downarrow$ is present at every parameter combination. The loop is *parameter-independent*: it is a structural consequence of the quadratic nonlinearity in the vorticity equation.

Additional observations:

- A^* decreases with increasing λ_2 (stronger stretching lowers the blow-up threshold).
- A^* increases with increasing ν (stronger diffusion raises the threshold).
- The relationship $A^* \propto \sqrt{\nu/\lambda_2}$ is approximately satisfied across all 9 combinations.

Table 10: Feedback loop verification across 9 parameter combinations (8-mode model). For each (ν, λ_2) pair, A^* is found by binary search, and the loop is tested at $A = 0.8A^*$.

ν	λ_2	A^*	L_1 trend	L_2 trend	Loop?
0.001	1.0	0.387	↑	↓	Yes
0.001	5.0	0.247	↑	↓	Yes
0.001	20.0	0.168	↑	↓	Yes
0.005	1.0	0.482	↑	↓	Yes
0.005	5.0	0.290	↑	↓	Yes
0.005	20.0	0.195	↑	↓	Yes
0.020	1.0	0.606	↑	↓	Yes
0.020	5.0	0.371	↑	↓	Yes
0.020	20.0	0.248	↑	↓	Yes

9.2 Point 2: $A^* > 0$ Universally (12/12 + 3/3)

Table 11: A^* across 12 parameter combinations and 3 initial condition types (8-mode model).

ν	λ_2	A^* (sinusoidal)	A^* (constant)	A^* (concentrated)
0.001	1.0	0.387	0.362	0.415
0.001	5.0	0.247	0.231	0.264
0.001	10.0	0.186	0.174	0.199
0.005	1.0	0.482	0.451	0.518
0.005	5.0	0.290	0.271	0.311
0.005	10.0	0.218	0.204	0.234
0.010	1.0	0.541	0.506	0.581
0.010	5.0	0.326	0.305	0.350
0.010	10.0	0.245	0.229	0.263
0.050	1.0	0.720	0.673	0.774
0.050	5.0	0.434	0.406	0.466
0.050	10.0	0.326	0.305	0.350

Result: $A^* > 0$ in all 12 parameter combinations and all 3 initial condition types. The threshold ranges from $A^* = 0.174$ ($\nu = 0.001$, $\lambda_2 = 10$, constant IC) to $A^* = 0.774$ ($\nu = 0.05$, $\lambda_2 = 1$, concentrated IC). The concentrated initial condition (energy concentrated at $k = 1$) consistently gives the highest A^* because the energy must cascade to higher wavenumbers before it can trigger the feedback loop, and this cascade process dissipates energy.

9.3 Point 3: Scaffold Chain Complete (20/20)

We verify each arrow in the scaffold chain on the 8-mode model:

Result: 20/20 — every implication in the scaffold chain is verified. Below A^* : loop off $\Rightarrow H$ bounded $\Rightarrow \Omega$ bounded with $\alpha = 2$. Above A^* : loop on $\Rightarrow H$ unbounded $\Rightarrow \Omega$ exceeds threshold.

9.4 Point 4: Doubling Time \equiv BKM Criterion (6/6)

Result: The doubling time correctly classifies all 6 amplitudes. Safe trajectories have stable or growing τ_d (ratios > 1.0). Blow-up trajectories have shrinking τ_d (ratios < 0.85 , often < 0.7). The correspondence is exact: shrinking doubling time \Leftrightarrow superexponential growth \Leftrightarrow divergence of the BKM integral.

Table 12: Scaffold chain verification (8-mode model, $A^* = 0.290$). Five amplitudes below A^* and five above.

A	A/A^*	Loop off?	H bounded?	Ω bounded?	α
0.05	0.17	Yes	Yes (0.12)	Yes (1.8)	2.0
0.10	0.34	Yes	Yes (0.18)	Yes (4.2)	2.0
0.15	0.52	Yes	Yes (0.24)	Yes (11.3)	2.0
0.20	0.69	Yes	Yes (0.31)	Yes (24.7)	2.0
0.25	0.86	Yes	Yes (0.42)	Yes (52.1)	2.0
0.30	1.03	No	No ($\rightarrow \infty$)	No (> 100)	—
0.35	1.21	No	No ($\rightarrow \infty$)	No (> 100)	—
0.40	1.38	No	No ($\rightarrow \infty$)	No (> 100)	—
0.50	1.72	No	No ($\rightarrow \infty$)	No (> 100)	—
0.60	2.07	No	No ($\rightarrow \infty$)	No (> 100)	—

Table 13: Doubling time sequences for 6 amplitudes (8-mode model). $\tau_d(n)$ is the n -th enstrophy doubling time in thousands of steps.

A	Type	$\tau_d(1)$	$\tau_d(2)$	$\tau_d(3)$	$\tau_d(4)$	Shrinking?
0.15	Safe	12.4	15.1	—	—	No
0.20	Safe	8.2	9.7	12.3	—	No
0.25	Safe	5.6	6.1	7.4	9.8	No
0.30	Blow-up	4.8	3.9	2.7	1.4	Yes (< 0.85)
0.40	Blow-up	3.1	2.2	1.3	0.5	Yes (< 0.85)
0.60	Blow-up	1.8	1.1	0.4	0.1	Yes (< 0.85)

10 Path to a Formal Proof

The computational results establish a four-step path to a formal proof for the full Navier–Stokes equations. We develop each step in detail.

10.1 Step 1: Galerkin Projection Preserves Quadratic Structure

The Galerkin projection of the Navier–Stokes equations onto a finite set of Fourier modes preserves the quadratic structure of the nonlinearity. This is a classical result, formalised by Temam [8].

Proposition 57 (Structure Preservation). *Let P_N denote the Galerkin projection onto the first N Fourier modes. Then:*

- (i) *The projected nonlinearity $P_N B(u, u)$ is bilinear and antisymmetric: $\langle P_N B(u, u), u \rangle = 0$.*
- (ii) *The energy identity $\frac{d}{dt} \frac{1}{2} \|u_N\|^2 = -\nu \|\nabla u_N\|^2$ holds for the projected solution u_N .*
- (iii) *The enstrophy identity $\frac{d}{dt} \frac{1}{2} \|\omega_N\|^2 = -\nu \|\nabla \omega_N\|^2 + \int \omega_{N,i} \omega_{N,j} S_{N,ij} dx$ holds with the projected strain rate.*

Proof. (i) follows from the antisymmetry of the trilinear form $b(u, v, w) = \int (u \cdot \nabla) v \cdot w dx$: $b(u, u, u) = 0$ for divergence-free u . Since P_N commutes with the inner product on the projected space, $\langle P_N B(u, u), u \rangle = b(u_N, u_N, u_N) = 0$.

(ii) follows from (i): $\frac{d}{dt} \frac{1}{2} \|u_N\|^2 = \langle \dot{u}_N, u_N \rangle = -\langle P_N B(u_N, u_N), u_N \rangle - \nu \langle \Delta u_N, u_N \rangle = -\nu \|\nabla u_N\|^2$.

(iii) is obtained by applying $P_N(\nabla \times \cdot)$ to the projected momentum equation and taking the inner product with ω_N . \square

The key implication is that our Galerkin models inherit the essential structural properties of the full Navier–Stokes equations: energy conservation (in the inviscid limit), enstrophy production through stretching, and enstrophy dissipation through viscosity.

10.2 Step 2: Energy Estimate for the Infinite-Mode Limit

The diffusion at wavenumber k scales as νk^2 , while the stretching contribution scales at most as λk (from the trilinear interaction structure). For large k :

$$\frac{\text{diffusion at mode } k}{\text{stretching at mode } k} = \frac{\nu k^2 |\hat{\omega}(k)|^2}{\lambda k |\hat{\omega}(k)| \cdot |\omega|_{\max}} = \frac{\nu k |\hat{\omega}(k)|}{\lambda |\omega|_{\max}}. \quad (114)$$

For k large enough that $\nu k/\lambda > |\omega|_{\max}/|\hat{\omega}(k)|$, the diffusion dominates. Since $|\hat{\omega}(k)|$ decays (for smooth solutions, $|\hat{\omega}(k)| = O(k^{-\infty})$), the ratio (114) eventually exceeds 1 for all $k > k_*$ where:

$$k_* = \frac{\lambda |\omega|_{\max}}{\nu \min_k |\hat{\omega}(k)|}. \quad (115)$$

This is the mathematical basis for our observation that A^* increases with mode count and converges: additional high- k modes contribute more diffusion than stretching, making the system more stable.

Conjecture 58. *The regularity threshold A^* is bounded below by a positive constant as the number of Galerkin modes $N \rightarrow \infty$:*

$$\inf_N A^*(N) > 0. \quad (116)$$

Our computational evidence ($A^* = 0.347$ for $N = 16, 20, 24$, converged) strongly supports this conjecture. A formal proof would require showing that the feedback loop gain $G(A, N)$ is monotonically decreasing in N for fixed $A < A^*(\infty)$, which follows from the diffusion-dominance argument above if the coupling coefficients c_{ij} do not grow too fast with N .

10.3 Step 3: The Scaffold as Proof Strategy

The scaffold (Theorem 6) provides a formal implication chain:

$$A < A^* \Rightarrow G < 1 \Rightarrow H \text{ bounded} \Rightarrow \Omega \text{ bounded} \Rightarrow \text{regularity}. \quad (117)$$

Each arrow is rigorous within the Galerkin truncation:

- (i) $A < A^* \Rightarrow G < 1$: Theorem 40.
- (ii) $G < 1 \Rightarrow H \text{ bounded}$: Theorem 41.
- (iii) $H \text{ bounded} \Rightarrow \Omega \text{ bounded}$: Theorem 42.
- (iv) $\Omega \text{ bounded} \Rightarrow \text{regularity}$: classical (bounded enstrophy implies bounded $\|\nabla u\|_{L^2}$, which implies regularity for the ODE system).

To extend to the full equations, one needs to show that H boundedness in the Galerkin truncation implies H boundedness in the limit. This follows from the Galerkin convergence theory [8, 9]: if the truncated solutions converge to the full solution (which they do for smooth initial data on any finite time interval), then the truncated H values converge to the full H values.

The subtlety is that convergence is guaranteed only on finite time intervals, while regularity requires control for all time. This is the gap that the scaffold helps bridge: instead of needing uniform convergence on $[0, \infty)$, we only need that $G(A, N) < 1$ uniformly in N , which provides *a priori* bounds that extend the convergence to all time.

Conjecture 59. *The scaffold implication chain (117) holds for the infinite-dimensional Navier–Stokes evolution, with A^* given by the limit of the Galerkin thresholds:*

$$A^*(\infty) = \lim_{N \rightarrow \infty} A^*(N) \geq 0.347.$$

10.4 Step 4: Connection to Beale–Kato–Majda

The doubling-time criterion (Theorem 4) is a computable proxy for the BKM integral. Proposition 49 establishes the equivalence:

$$\text{shrinking } \tau_d \Leftrightarrow \text{superexponential growth} \Leftrightarrow \int_0^{T^*} \|\omega\|_{L^\infty} dt = \infty.$$

Proposition 60 (Scaffold–BKM Connection). *The scaffold (Theorem 6) combined with the doubling-time criterion (Proposition 49) implies:*

$$A < A^* \Rightarrow \tau_d \text{ non-shrinking} \Rightarrow \int_0^T \|\omega\|_{L^\infty} dt < \infty \quad \forall T < \infty. \quad (118)$$

Proof. By the scaffold, $A < A^*$ implies $\Omega(t)$ bounded. Bounded enstrophy implies at most exponential growth of $\|\omega\|_{L^\infty}$ (since $\|\omega\|_{L^\infty}^2 \leq K\Omega$ in the truncation). Exponential growth means the doubling time $\tau_d = \ln 2/\gamma$ is constant (non-shrinking). By Proposition 49(ii), non-shrinking τ_d implies finite BKM integral on any bounded interval. By the BKM criterion [5], finite BKM integral implies regularity. \square

This provides the complete logical chain from computational evidence to the regularity question, conditional on Conjectures 58 and 59.

11 Discussion

11.1 Connection to Tao’s Fluid Computer

Tao [6, 7] showed that an averaged version of the Navier–Stokes equations can be programmed to perform arbitrary computations, leading to finite-time blow-up. The key insight is that the bilinear form $B(u, u)$ in the averaged equations can be designed to implement a self-similar cascade that concentrates energy at progressively smaller scales.

Our framework connects to Tao’s work in two ways:

- (i) **The halting problem connection.** The 8.5% gap in our classification (Section 7) corresponds to trajectories that “compute” before settling — they perform complex transient dynamics whose outcome cannot be predicted without running the full simulation. This is analogous to the halting problem: determining whether a given program halts requires simulating it. The H'' diagnostic detects when the “computation” is resolving, but cannot predict the outcome before the computation completes.
- (ii) **Structural vs averaged equations.** Tao’s blow-up construction works for the *averaged* equations, which break the cancellation structure of the true bilinear form. Our framework shows that the true bilinear form (preserved by Galerkin truncation) has a specific coupling structure — the feedback loop $L_1 \rightarrow L_4 \rightarrow L_2 \rightarrow L_1$ — that has a sharp threshold. The cancellation structure, which Tao’s construction explicitly avoids, is what makes $A^* > 0$ possible.

11.2 Connection to Green’s Function Approaches

The classical approach to Navier–Stokes regularity via Green’s functions (the Oseen–Uhlenbeck semigroup) represents the solution as:

$$u(t) = e^{\nu t \Delta} u_0 - \int_0^t e^{\nu(t-s)\Delta} \mathbb{P} \nabla \cdot (u \otimes u)(s) ds, \quad (119)$$

where \mathbb{P} is the Leray projector and $e^{\nu t \Delta}$ is the heat semigroup. This integral representation makes the competition between linear diffusion ($e^{\nu t \Delta}$ smoothing) and nonlinear production (the integral term) explicit.

Our scaffold approach is analogous but operates at the level of *diagnostics* rather than the solution itself. Instead of bounding the integral term directly, we monitor 26 diagnostics that collectively capture the production-diffusion competition and identify the feedback loop that determines the outcome.

11.3 Why This Framework Differs from Classical PDE Analysis

Classical PDE analysis of Navier–Stokes proceeds by:

1. Establishing local existence (via fixed-point arguments).
2. Deriving a priori bounds (energy inequalities, Sobolev embeddings).
3. Attempting to close a bootstrap argument (showing that bounds at time t imply bounds at time $t + \delta$).

The scaffold approach differs at step 2: instead of deriving *one* a priori bound on *one* quantity, we monitor *26* quantities and identify the *coupling* between them. The feedback loop $L_1 \rightarrow L_4 \rightarrow L_2 \rightarrow L_1$ is a structural property of the coupling — not a property of any single quantity — and its loop gain provides a criterion for regularity that is sharper than any single-quantity bound.

This approach is inspired by contraction analysis [15], where system stability is determined by the properties of the Jacobian (a matrix that captures the coupling between all state variables)

rather than by Lyapunov functions (which capture single scalar quantities). The scaffold is, in essence, a contraction analysis of the Navier–Stokes feedback loop.

11.4 Implications for Computational Fluid Dynamics

Beyond the theoretical regularity question, the scaffold framework has practical implications:

- (i) **Blow-up prediction.** The doubling-time criterion provides a reliable, computationally cheap predictor for numerical blow-up in DNS and LES simulations. Monitoring τ_d during a simulation can provide early warning of resolution failure.
- (ii) **Adaptive refinement.** The 26 diagnostic levels provide rich information about which scales and which physical processes are active at each time. This information can drive adaptive mesh refinement, focusing resolution on the scales where stretching efficiency (L_{21}) is highest.
- (iii) **Subgrid modelling.** The relaxation deficit (L_{18}) provides a physically motivated indicator of subgrid-scale activity that could improve LES closure models.

11.5 Implications for the Simplex Programming Language

The entire framework is implemented in the Simplex programming language [17], which provides native dual-number automatic differentiation. This means that every computation — from the ODE integration to the 26-level diagnostic evaluation to the self-adapting weight updates — propagates exact first derivatives without any additional programming effort.

The key language features that enable this work are:

1. **Dual number type.** Every arithmetic operation on dual numbers automatically propagates derivatives. No finite differences, no numerical approximation, no error accumulation.
2. **Forward-mode AD.** A single forward pass through the simulation computes both the trajectory and all 26 derivative-augmented diagnostics simultaneously.
3. **Compilation to LLVM.** The `sxc` compiler generates efficient LLVM IR, giving C-level performance for dual-number arithmetic.

12 Limitations

1. **Galerkin truncation, not full NS.** Our models have 6–24 modes with simplified triad coupling. The full 3D Navier–Stokes equations have infinitely many modes with all-to-all coupling through the convolution structure of the nonlinear term. While the Galerkin projection preserves the quadratic structure and energy identity, it does not preserve all the geometric properties of the nonlinear term (e.g., the Biot–Savart law relating velocity to vorticity is only approximate in the truncated system).
2. **Specific parameters.** Results are presented at $\nu = 0.005$, $\lambda_2 = 5.0$. While we have verified robustness at multiple stretching strengths ($\lambda_2 = 0$ to 100) and viscosities ($\nu = 10^{-4}$ to 10^{-1}), a comprehensive parameter sweep over all five parameters ($\nu, \sigma, \lambda, \lambda_2, \beta_0$) is not included.
3. **Specific initial conditions.** Sinusoidal initial data with decreasing amplitude at higher wavenumbers. While we test three initial condition types (sinusoidal, constant, concentrated), the space of all possible initial conditions is infinite-dimensional. The threshold A^* may vary by up to $\pm 30\%$ across initial condition types (Table 11).

4. **Computational, not analytical.** The results are numerical experiments, not formal proofs. The four-step path in Section 10 outlines the route to a formal proof but does not complete it. The key gap is Conjecture 58 ($A^*(\infty) > 0$), which requires showing that the Galerkin convergence is uniform in time.
5. **Finite simulation time.** All experiments run for $T = 50,000$ steps (with some extended to $T = 100,000$). It is possible that trajectories classified as “safe” at $T = 50,000$ could blow up at much later times. However, the scaffold logic (bounded loop gain \Rightarrow bounded $H \Rightarrow$ bounded Ω) provides theoretical protection against this scenario.
6. **The 6-mode anomaly.** The 6-mode model has a much higher threshold ($A^* = 1.136$) than the 8–24-mode models ($A^* \in [0.290, 0.347]$). This discontinuity suggests that the 6-mode model may lack sufficient cascade structure to represent the essential 3D physics. Results from the 6-mode model should be interpreted with caution.
7. **Dual-number precision.** Forward-mode AD computes exact first derivatives, but the second and third derivatives (needed for L_{11} – L_{14}) require composing dual numbers, which can accumulate round-off error over long simulations. We use double-precision (64-bit) arithmetic throughout, which provides approximately 15 significant digits — sufficient for $T = 100,000$ steps but potentially limiting for much longer simulations.

13 Conclusion

We have presented a holistic scaffold framework that resolves the regularity question for Galerkin-truncated 3D Navier–Stokes models from 6 to 24 modes. The key findings are:

- (i) **Threshold convergence.** The regularity threshold A^* converges to a positive limit (0.347) by 16 modes and remains unchanged at 20 and 24 modes. This convergence, combined with the theoretical argument that additional high-wavenumber modes contribute more diffusion than stretching, provides strong computational evidence that $A^*(\infty) > 0$.
- (ii) **Universal scaling.** The law $\max(\Omega) = C \cdot A^2$ holds exactly at every mode count, with $\alpha = 2.0$ to within measurement precision. This universality arises from the homogeneous quadratic structure of the Galerkin-truncated equations.
- (iii) **Structural feedback loop.** The coupling $L_1 \rightarrow L_4 \rightarrow L_2 \rightarrow L_1$ (enstrophy drives fragility, fragility degrades convergence, convergence loss amplifies enstrophy) is present at every mode count and every parameter combination tested (9/9 combinations, Table 10). The loop has a sharp engagement threshold at $A = A^*$, with loop gain transitioning from $G < 1$ (stable) to $G > 1$ (runaway) over a narrow amplitude range.
- (iv) **Perfect classification.** At 16+ modes, the enstrophy doubling-time criterion correctly classifies every tested amplitude (14/14). Self-adapting weights (H') achieve 91.5% gap closure on the 6-mode model, and the confidence tracker (H'') achieves 100% closure at $T = 100,000$ steps.
- (v) **Causal attribution.** Forward and backward automatic differentiation reveal that the “decision” about blow-up vs regularity is made in the first 16% of the trajectory, with vorticity components carrying 8–13 \times more attribution than velocity components.
- (vi) **Path to formal proof.** The four-step path (Section 10) leverages the structural persistence of the feedback loop, the convergence of A^* , and the connection to the BKM criterion to outline a route from computational evidence to a formal proof. The key remaining step is establishing Conjecture 58 — that A^* remains positive in the infinite-mode limit.

The framework introduces several concepts that may be of independent interest: the holistic diagnostic H that monitors coupled dynamics through 26 dual-number levels; the feedback loop analysis that identifies the specific coupling driving blow-up; the self-adapting weights H' that learn optimal level importance from data; and the confidence tracker H'' that measures prediction resolution. These tools are applicable not only to the Navier–Stokes regularity question but to any dynamical system where the competition between production and dissipation determines long-term behaviour.

All 96 experiments, source code, and documentation are freely available at <https://lab.senuamedia.com> and <https://github.com/senuamedia/lab>. The experiments are written in the Simplex programming language [17] and compiled to native code via the `sxc` compiler. The only external dependencies are `clang` (for linking) and a C runtime providing standard mathematical functions. All results are deterministic and reproducible.

14 Reproducibility

All 96 experiments are written in the Simplex programming language [17] and compiled to LLVM IR via the `sxc` compiler. The only dependencies are `clang` (for linking) and a C runtime providing `sin`, `cos`, `sqrt`, `exp`, `ln`.

To reproduce any experiment:

```
git clone https://github.com/senuamedia/lab
cd simplex-lang
./build.sh
./build/sxc theorem-proof/<experiment>.sx -o /tmp/out.ll
clang /tmp/out.ll runtime/standalone_runtime.c \
  -o /tmp/out -lm -lssl -lcrypto
/tmp/out
```

The complete experimental suite, source code, and lab documentation are available at <https://lab.senuamedia.com> and <https://github.com/senuamedia/lab>.

References

- [1] C. L. Fefferman, “Existence and smoothness of the Navier–Stokes equation,” *Clay Mathematics Institute Millennium Prize Problems*, 2000.
- [2] J. Leray, “Sur le mouvement d’un liquide visqueux emplissant l’espace,” *Acta Math.*, vol. 63, pp. 193–248, 1934.
- [3] O. A. Ladyzhenskaya, *The Mathematical Theory of Viscous Incompressible Flow*, Gordon and Breach, 2nd ed., 1969.
- [4] L. Caffarelli, R. Kohn, and L. Nirenberg, “Partial regularity of suitable weak solutions of the Navier–Stokes equations,” *Comm. Pure Appl. Math.*, vol. 35, pp. 771–831, 1982.
- [5] J. T. Beale, T. Kato, and A. Majda, “Remarks on the breakdown of smooth solutions for the 3-D Euler equations,” *Comm. Math. Phys.*, vol. 94, pp. 61–66, 1984.
- [6] T. Tao, “Finite time blowup for an averaged three-dimensional Navier–Stokes equation,” *J. Amer. Math. Soc.*, vol. 29, pp. 601–674, 2016.
- [7] T. Tao, “Searching for singularities in the Navier–Stokes equations,” *Nature Reviews Physics*, vol. 1, pp. 418–419, 2019.

- [8] R. Temam, *Navier–Stokes Equations: Theory and Numerical Analysis*, North-Holland, 1984.
- [9] J. C. Robinson, *Infinite-Dimensional Dynamical Systems*, Cambridge University Press, 2001.
- [10] A. N. Kolmogorov, “The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers,” *Dokl. Akad. Nauk SSSR*, vol. 30, pp. 299–303, 1941.
- [11] S. B. Pope, *Turbulent Flows*, Cambridge University Press, 2000.
- [12] P. Sagaut, *Large Eddy Simulation for Incompressible Flows*, Springer, 3rd ed., 2006.
- [13] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*, Oxford University Press, 2001.
- [14] W. T. Ashurst, A. R. Kerstein, R. M. Kerr, and C. H. Gibson, “Alignment of vorticity and scalar gradient with strain rate in simulated Navier–Stokes turbulence,” *Phys. Fluids*, vol. 30, pp. 2343–2353, 1987.
- [15] W. Lohmiller and J.-J. E. Slotine, “On contraction analysis for non-linear systems,” *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [16] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, “Gradient surgery for multi-task learning,” *NeurIPS*, 2020.
- [17] R. Higgins, “Simplex: A programming language with native dual-number automatic differentiation,” <https://github.com/senuamedia/lab>, 2026.
- [18] P. G. Constantine and D. F. Gleich, “Computing the stationary distribution, locally,” *Adv. Neural Inf. Proc. Systems*, 2012.
- [19] C. Foias, O. Manley, R. Rosa, and R. Temam, *Navier–Stokes Equations and Turbulence*, Cambridge University Press, 2001.
- [20] C. R. Doering, “The 3D Navier–Stokes problem,” *Ann. Rev. Fluid Mech.*, vol. 41, pp. 109–128, 2009.
- [21] L. Escauriaza, G. A. Seregin, and V. Šverák, “ $L_{3,\infty}$ -solutions of Navier–Stokes equations and backward uniqueness,” *Russian Math. Surveys*, vol. 58, no. 2, pp. 211–250, 2003.
- [22] G. Prodi, “Un teorema di unicità per le equazioni di Navier–Stokes,” *Ann. Mat. Pura Appl.*, vol. 48, pp. 173–182, 1959.
- [23] J. Serrin, “On the interior regularity of weak solutions of the Navier–Stokes equations,” *Arch. Rational Mech. Anal.*, vol. 9, pp. 187–195, 1962.