

Continuum Blow-Up, Finite Distinguishability, and Physical Admissibility

A consolidated companion to the VERSF Navier–Stokes programme

Keith Taylor

15 September 2026 · Consolidated working manuscript 2.0

VERSF Theoretical Physics Programme

General reader's summary

The Navier–Stokes equations describe how water and air move. They treat a fluid as a continuous material, rather than following its individual molecules. A famous mathematical question asks whether these equations can start with a smooth flow and later predict a speed that grows without limit. Such a breakdown is called a **singularity**. [9]

Our original VERSF paper explored how this might happen. It also asked whether the ever-smaller details in the equations could all represent real physical differences. It did not solve the mathematical problem, and a later review found errors in parts of its argument. Those errors are recorded here; the results below do not rely on them. [1; Appendix A]

OpenAI subsequently claimed a mathematical solution: a carefully chosen force drives a fluid from rest towards an unbounded speed, even though its total energy stays bounded. A correct proof of that kind can meet the mathematical challenge. It is not an experiment showing real water or air reaching infinite speed. [6, 9]

This companion examines two different questions. **First, what would particular measurements report?** The speed at one tiny spot is not the same as the average motion across a patch. We show that, under stated conditions, measurements of the latter kind can settle to finite answers while the calculated peak keeps growing. The measurements must keep the same limits on sensitivity and detail; we do not establish this for every possible experiment.

Second, how long can these equations still describe a real fluid accurately? For a chosen material, there may be a smallest scale at which treating it as continuous is reliable. There may also be a speed range within which ignoring compression is justified. Using the construction's scale laws, we calculate when specified limits of those kinds must be crossed—before the formal singular moment. These calculations depend on the material assumptions; they do not impose one smallest length on all of nature.

Together, these results explain why an infinity in the equations is not, by itself, proof of an infinity in a real fluid. They do not disprove a correct OpenAI proof, and they do not tell us what replaces the predicted flow after the model stops being reliable.

VERSF asks one further question: which differences can leave separate, lasting traces in nature, whether anyone watches or not? We still need to connect the fluid's interactions with its surroundings to that question. **The equations, the measurements, and the physical traces are related—but they are not the same thing.**

Abstract

This consolidated companion separates continuum blow-up, finite-resource readout behaviour, and the validity of a chosen material-fluid approximation. Its analytic core retains the self-contained finite-energy estimates of version 1.1. Fixed-resolution velocity and vector-mollified-vorticity readouts remain bounded; a solenoidal weak momentum identity yields a unique weak L^2 terminal state, uniform endpoint convergence, and finite total variation for probe families with specified common norm bounds. A vanishing-energy component is uniformly negligible to bounded linear probes. Finite-rank or observation-stable fluid-to-closure maps transfer the appropriate results into the existing VERSF operational metric, conditionally on physically identifying those maps. The quantitative addition derives a carrier wavelength law $\lambda_p \asymp Q^{(1+h)/2}$ from the cited construction's dyadic variables and proves a finite-band bound for a specified positive material validity scale. Separate, explicitly conditional Knudsen and low-Mach estimates locate exits from a stipulated hydrodynamic regime. Uniform bounds, dimensional calibration, variable material parameters, and the distinction between an equality-scale estimate and a guaranteed exit bound are kept explicit. The attachment's adjacent-refinement criterion is not adopted: small successive increments neither imply convergence nor finite total variation. The original mathematical audit and Operational Geometry correction record remain in force. No global regularity theorem, universal invisibility result, BCB verdict, physical impossibility theorem, or independent verification of the external construction is claimed.

Status and scope. AI-assisted working research manuscript, not independently peer-reviewed or machine-checked. This edition consolidates the first companion [15] with the quantitative material-model discussion [10]. It replaces neither the original VERSF manuscript nor the Operational Geometry source. It does not revise the separate explanatory commentary. New derivations and departures from the supplied drafts are recorded in Appendix C.

Keywords: Navier–Stokes; finite-time blow-up; finite distinguishability; weak terminal state; finite-resource readouts; continuum validity; pulse wavelength; Knudsen number; low-Mach approximation; VERSF.

1. The original question, the external claim, and this companion

1.1 What the original VERSF paper proposed

A Rigorous Reduction of Potential Finite-Time Breakdown in 3D Navier–Stokes [1] studied the unforced equations. It organised a proposed vortex-growth route around sustained self-stretching, velocity-gradient control, and limited cancellation, labelled A, B, and C. It did not establish that any specified initial flow maintains the proposed conditions. Its physical discussion also asked whether a formal continuum solution must represent a physically admissible fluid, and introduced a BCB scale-budget interpretation.

The later audit found errors inside several estimates and the conditional Riccati argument, not merely unproved persistence hypotheses. This companion preserves that audit in Appendix A and does not reinstate the disputed theorems. The original paper supplies the research question and proposed mechanism; Sections 3–5 below supply independent analysis.

1.2 What OpenAI claims—and what the Clay question requires

Theorem 1.1 of OpenAI’s *Finite Time Blowup for Navier–Stokes* [6] states a solution starting from rest, smooth before $T = 1$, with a smooth force compactly supported in spacetime and uniformly bounded kinetic energy, but unbounded maximum velocity near T . The manuscript identifies this with the permitted forced breakdown cases C and D of the official Clay formulation [9]. We use its stated properties as source inputs; we have not independently verified its complete proof or Lean formalisation.

The official question permits suitably regular external forcing in its breakdown alternatives. Thus allowed forcing does not make a construction conditional on an unproved conjecture, and a correct forced example need not settle the unforced case. Mathematical existence is also not a laboratory demonstration of infinite speed. [9, alternatives A–D]

The announcement is dated 8 September 2026 [7]. The CMI page and signed statement both display 11 September 2026 [8]; they describe an apparent settlement while reserving evaluation. No later acceptance decision is inferred here. OpenAI’s announcement itself recognises the distinction between a singular continuum model and a material fluid.

1.3 What the first companion establishes

The first companion, version 1.1 [15], changed the object under examination: it studied bounded families of readouts rather than trying to prove or exclude pointwise blow-up. Sections 3–8 and Appendices A–B of this edition retain that analysis and its qualifications. A finite weak endpoint is not asserted to be smooth, and smoothing a nonlinear solution does not create another solution with the same data and force.

The central mathematical pairing is **Theorem 5.1 with Section 8.1**. The theorem gives convergence and finite variation for every family satisfying common probe bounds. The counterexample shows why saying that each individual probe is finite cannot supply a bound for the whole family. Whether physically admissible record-forming interactions possess such a common budget remains an identifiable physical question, not an assumption hidden in the word “finite.”

1.4 What the quantitative attachment adds

The supplied *Continuum Blow-Up and Physical Admissibility* draft [10] studies the shrinking pulse scales and material-model limits. Sections 9–10 integrate those calculations, with explicit constants and physical assumptions. Section 8.6 replaces its adjacent-stage distinguishability test rather than adopting it as an additional VERSF axiom.

The resulting structure is deliberate. Sections 2–8 concern distinguishability and readouts; Sections 9–10 concern the scale hierarchy and the validity of a specified material approximation; Section 11 states the physical connection still needed. The former mathematical results do not depend on a molecular cutoff, and the latter conditional cutoff results do not depend on BCB or on VERSF being correct. Their combination concerns the interpretation of a claimed solution, not a competing solution to the Clay problem.

2. What the VERSF sources actually require

2.1 Finite distinguishability is not BCB

Section 8.2 of *On the Assumption of Distinguishability in Physical Theory* [2] states:

“Any finite region of the substrate supports a finite number of distinguishable configurations.”

Its companion conditions concern irreversible commitment and finite local capacity. Section 9.7 explicitly allows continuous mathematical descriptions while declining to identify every mathematical difference with a separately realisable Fact. A distinction must be capable of entering a stable record through a finite physical interaction. This is an interaction-based criterion, not a requirement for a human observer.

Facts as Structural Necessities [3, Definition 2] describes a distinguishability structure d and a positive separation threshold $\delta_{\min}$ for operationally separable alternatives. We use this as part of the VERSF framework. We do not claim that finite-energy fluid analysis proves the foundational derivation offered in [3].

By contrast, Appendix B of [1] proposes the scale-budget axiom

$$I_{\text{stretch},\ell} \leq \vartheta E_\ell + C_0 M_\ell + R_\ell, \quad 0 < \vartheta < 1, \quad R_\ell = o(M_\ell^2).$$

Here $M_\ell = \|K_\ell * \omega\|_\infty$, and E_ℓ denotes that manuscript’s intervention bound. This is BCB as used in the Navier–Stokes manuscript. It constrains a proposed dynamical budget; it is not the finite-distinguishability axiom. Moreover, E_ℓ is built from error bounds. Its magnitude cannot automatically be interpreted as the amount of signed physical dissipation. Section B.4 of [1] correctly cautions that a lower growth inequality cannot supply an upper growth bound, and that its argument does not exclude all blow-up mechanisms.

2.2 The operational metric already exists in the corpus

Operational Geometry in the VERSF Framework [5, Definitions 3.4–3.7 and Theorem 4.1] introduces an operational subspace and metric of the form

$$A = \text{Im}(\Omega_{\max}), \quad D_{\text{op}}(a, b) = \|a - b\|_{\text{op}}, \quad \dim A = d < \infty.$$

The paper distinguishes this continuous finite-dimensional space from the discrete admissible sector set Σ . Its projection is non-expansive in its **native closure-state norm** [5, Theorem 5.1]. That is not automatically a bound in the fluid L^2 norm.

Accordingly, the missing object is not another distinguishability axiom or another abstract metric. It is a physically justified map from fluid configurations and their record-forming environment into the existing closure geometry. Section 7 states exactly what properties of that map suffice.

2.3 No universal spatial cutoff is assumed

An operational separation $\delta_{\min}$ is not, without a calibration, a length in metres. A fixed mathematical averaging scale q is likewise not automatically the closure scale ξ .

The proto-factual manuscript [4, Section 3] explicitly permits sub-coherence processes to carry energy, momentum, and causal influence without each becoming an independent committed Fact. It also describes ξ as a causal-capacity threshold, not simply a contiguous geometric size. We therefore do not infer that all sub- ξ activity is forbidden, or that Navier–Stokes must stop at a prescribed universal wavelength.

A further distinction is essential: a continuous family of mathematical configurations need not encode infinitely many simultaneously committed records. Neither the number of continuum points nor the number of refinements in a mathematical description is, by itself, a count of Facts.

3. Analytic setting and notation

Let $0 < T < \infty$, $\nu > 0$, and consider

$$\partial_t u + \operatorname{div}(u \otimes u) - \nu \Delta_x u + \nabla p = f, \quad \operatorname{div} u = 0$$

for $x \in \mathbb{R}^3$ and $0 \leq t < T$. Write $\omega = \nabla \times u$. Spatial norms use Lebesgue measure. The mathematical kinetic energy is $\frac{1}{2} \|u\|_2^2$; a fixed physical mass-density factor does not change the estimates.

We assume that the solution is smooth before T , has sufficient decay for the usual solenoidal weak formulation, and satisfies

$$U := \sup_{0 \leq t < T} \|u(t)\|_2 < \infty, \quad f \in L^1(0, T; L^2).$$

More precisely, our dynamical results use the identity

$$\frac{d}{dt} \langle u, \phi \rangle = \langle u \otimes u, \nabla \phi \rangle + \nu \langle u, \Delta_x \phi \rangle + \langle f, \phi \rangle$$

for time-independent solenoidal test fields ϕ . Compactly supported smooth velocity and pressure, as in the cited construction, justify the integrations directly. The results also apply whenever this weak identity is valid by approximation for the test classes specified below. This formulation avoids imposing an unstated condition on an arbitrary harmonic pressure component.

Let K be a nonnegative Schwartz kernel with $\int K = 1$; a Gaussian is a convenient choice. Set

$$K_\varrho(x) = \varrho^{-3} K(x/\varrho), \quad S_\varrho u = K_\varrho * u, \quad \varrho > 0.$$

The resolution is ϱ ; the spatial Laplacian remains Δ_x . For a multi-index α , put $m = |\alpha|$. Then

$$\|D^\alpha K_\varrho\|_2 = \varrho^{-m-3/2} \|D^\alpha K\|_2.$$

Readouts and thresholds below are calibrated in a fixed set of units. Equivalently, each probe includes its reference output scale so its reported value is dimensionless. No physical identification is made by equating dimensionally different quantities.

4. Fixed-resolution boundedness and vanishing-energy concentration

Proposition 4.1. Energy controls every fixed-scale derivative readout

For $u \in L^2(\mathbb{R}^3; \mathbb{R}^3)$,

$$\|D^\alpha S_\varrho u\|_\infty \leq \|D^\alpha K\|_2 \varrho^{-m-3/2} \|u\|_2.$$

Proof. At each x , apply Cauchy–Schwarz to $\int D^\alpha K_\varrho(x - y) u(y) dy$, then take the supremum in x . Translation leaves the kernel norm unchanged. ▀

In particular, distributional differentiation commutes with convolution, so

$$\|K_\varrho * \omega\|_\infty = \|\nabla \times S_\varrho u\|_\infty \leq C_K \varrho^{-5/2} \|u\|_2.$$

Thus a uniformly finite-energy solution cannot have its vector-mollified vorticity diverge at any one fixed positive resolution. This directly corrects the fixed-scale claim in Section 0.3 of [1]. It does not exclude divergence when the resolution itself tends to zero.

The distinction is in the order and uniformity of the limits. At a smooth time, the approximate identity recovers ω as $\varrho \downarrow 0$. The constants in Proposition 4.1 are not uniform in that limit. Boundedness for each fixed ϱ therefore does not control the unsmoothed pointwise norm.

Proposition 4.2. A vanishing-energy component is uniformly small to bounded linear probes

Let $v(t) \in L^2$ satisfy $\|v(t)\|_2 \rightarrow 0$ as $t \uparrow T$. Let P be a family of vector probes with

$$B := \sup_{\phi \in P} \|\phi\|_2 < \infty.$$

Then

$$\sup_{\phi \in P} |\langle v(t), \phi \rangle| \leq B \|v(t)\|_2 \rightarrow 0.$$

Proof. Cauchy–Schwarz gives the bound for each probe, with the same constant B . Taking the supremum preserves that bound. ▀

The word **uniformly** is doing real work. The theorem applies to a whole bounded family, not merely to one detector at a time. It applies, for example, to all translated copies of a fixed-resolution averaging kernel and to their derivatives of any fixed order.

For a decomposition $u = b + v$, the conclusion is convergence of the selected readouts of $u - b$. It does not establish that b solves Navier–Stokes or has a smooth terminal limit. In particular, a cutoff decomposition is an analytic device, not a decomposition into two independent fluid solutions.

5. A terminal-state theorem for the complete flow

The preceding proposition concerns a selected small component. The weak momentum equation gives a separate result for the **entire** flow, without removing the core or ignoring oscillatory corrections.

5.1 A bounded class of dynamical readouts

Let Φ be a family of time-independent solenoidal test fields for which the weak identity in Section 3 holds. Assume

$$\sup_{\phi \in \Phi} \|\phi\|_2 \leq B_0, \quad \sup_{\phi \in \Phi} \|\nabla \phi\|_\infty \leq B_1, \quad \sup_{\phi \in \Phi} \|\Delta_x \phi\|_2 \leq B_2.$$

The gradient norm is the pointwise matrix norm compatible with contraction against $u \otimes u$. Define the readout pseudometric

$$d_\Phi(u, v) := \sup_{\phi \in \Phi} |\langle u - v, \phi \rangle|.$$

It is finite on L^2 , symmetric, and satisfies the triangle inequality. It may fail to distinguish some different fields, so it is called a pseudometric. This is a specified analytic readout model, not a claim that every physically admissible interaction belongs to Φ .

Theorem 5.1. Weak terminal state, uniform operational convergence, and finite variation

Under the assumptions of Section 3 there exists a unique $u_* \in L^2_\sigma$, with $\|u_*\|_2 \leq U$, such that

$$u(t) \rightharpoonup u_* \quad \text{weakly in } L^2 \text{ as } t \uparrow T.$$

For every probe family Φ as above,

$$d_\Phi(u(t), u_*) \leq (B_1 U^2 + \nu B_2 U)(T - t) + B_0 \int_t^T \|f(s)\|_2 ds.$$

Moreover, after assigning $u(T) = u_*$ in this pseudometric, the observable path has total variation bounded by

$$\text{Var}_\Phi(u; [0, T]) \leq L_\Phi := (B_1 U^2 + \nu B_2 U)T + B_0 \int_0^T \|f(s)\|_2 ds < \infty.$$

Proof. For each admissible ϕ , the weak equation and Hölder's inequality give

$$\left| \frac{d}{dt} \langle u, \phi \rangle \right| \leq \|\nabla \phi\|_\infty \|u\|_2^2 + \nu \|\Delta_x \phi\|_2 \|u\|_2 + \|\phi\|_2 \|f\|_2.$$

Thus every compactly supported smooth solenoidal test has a limiting pairing at T . The L^2 bound gives weakly convergent subsequences along any sequence of times approaching T . Any two subsequential limits have the same pairing with every such test. These tests are dense in L^2_σ , so the limits coincide. This establishes the unique weak limit u_* and its energy bound.

For $t < s < T$, integrate the displayed inequality and use the uniform bounds B_0, B_1, B_2 . The result is

$$d_\phi(u(t), u(s)) \leq (B_1 U^2 + \nu B_2 U)(s - t) + B_0 \int_t^s \|f(a)\|_2 da.$$

For any fixed $\phi \in \Phi$, weak convergence supplies its terminal pairing. Let $s \uparrow T$ in the bound for that pairing, then take the supremum over Φ . The bound is independent of ϕ , so this gives the claimed uniform terminal estimate without an invalid interchange of a supremum and a limit.

Finally, apply the two-time estimate to each interval in a finite partition of $[0, T]$ and add. The time lengths and forcing integrals telescope to the stated bound. Taking the supremum over partitions proves finite variation. ▀

Scope. The theorem does not assert that u_* is smooth, that $u(t) \rightarrow u_*$ strongly in L^2 , or that a classical continuation through T exists. A weak finite-energy endpoint and a pointwise singularity are compatible.

Corollary 5.2. Fixed-resolution velocity and vorticity have finite endpoint limits

For every fixed $q > 0$ and every finite multi-index α ,

$$D^\alpha S_q u(t) \rightarrow D^\alpha S_q u_* \quad \text{uniformly in space as } t \uparrow T.$$

An explicit bound is

$$\begin{aligned} \|D^\alpha S_q(u(t) - u_*)\|_\infty &\leq C_{K,m} (q^{-m-4} U^2 + \nu q^{-m-7/2} U)(T - t) \\ &+ C_{K,m} q^{-m-3/2} \int_t^T \|f(s)\|_2 ds. \end{aligned}$$

Proof. Let P be the L^2 Leray projection onto solenoidal fields. Since u is solenoidal, a component of $D^\alpha S_q u(x)$ is the pairing of u with a translated projected kernel

$$\phi_{x,j} = P[D^\alpha K_q(x - \cdot) e_j].$$

Projection is an L^2 contraction and commutes with constant-coefficient derivatives. Consequently, the B_0 and B_2 bounds are proportional to $q^{-m-3/2}$ and $q^{-m-7/2}$. The Leray Fourier multiplier has bounded matrix norm. Writing $\hat{K}$ for the Fourier transform of K , Fourier inversion gives

$$\|\nabla \phi_{x,j}\|_\infty \leq C \int_{\mathbb{R}^3} |k|^{m+1} |\hat{K}(qk)| dk \leq C_{K,m} q^{-m-4}.$$

These estimates are uniform in x and j . The projected kernels are generally not compactly supported and may have $|y|^{-3}$ tails, but every weak-form pairing is absolutely convergent: $\nabla \phi \in L^\infty$ pairs with $u \otimes u \in L^1$, while $\Delta_x \phi \in L^2$ and $\phi \in L^2$ pair with $u \in L^2$ and $f \in L^2$, respectively. For these kernels, a smooth Fourier cutoff away from zero followed by cutoff of the resulting Schwartz vector potential gives compactly supported solenoidal approximants converging in L^2 , in the L^2 norm of the Laplacian, and in the L^∞ norm of the gradient. Thus the integrated weak identity passes to the limit; for compactly supported u, p it also follows directly by integration by parts. Theorem 5.1 now applies. Taking the appropriate first-derivative combination gives the vorticity assertion. ▀

No hard spatial cutoff has been introduced. The theorem holds at **each chosen positive scale**, while its constants worsen as that scale is reduced.

Corollary 5.3. A bound on separated readout states along the trajectory

If $0 \leq t_1 < \dots < t_N \leq T$ and successive sampled fluid configurations satisfy

$$d_\phi(u(t_{i+1}), u(t_i)) \geq \delta > 0,$$

then

$$(N - 1)\delta \leq L_\phi, \quad N \leq 1 + \frac{L_\phi}{\delta}.$$

Proof. The sum of successive separations is bounded by the total variation from Theorem 5.1. ▫

This is a bound on distinguishable **fluid-state readouts in the specified class**. It is not a bound on every record produced by the environment, on external clock ticks, or on repeated recordings of the same state. Those records contain additional physical variables not represented by u alone.

6. Application to a concentrating core

The force in [6, Theorem 1.1, p. 1] belongs to $C_c^\infty(\mathbb{R}^3 \times (0, \infty); \mathbb{R}^3)$ and therefore to $L^1(0, 1; L^2)$; together with the theorem's bounded kinetic energy and compactly supported smooth pre-terminal velocity and pressure, this verifies Section 3's hypotheses and makes Theorem 5.1 and Corollary 5.2 applicable to that complete flow.

6.1 Source inputs and the limited role of the leading profile

For $\tau = 1 - t$, the core C_τ defined in [6, Section 3.1, p. 8] has the following size and velocity bounds:

$$\begin{aligned} \ell_r &= \tau^{1/2}, & \ell_z &= \tau^{1/2-h}, & 0 < h < 1/100, \\ \|u^{(0)}(1 - \tau)\|_{L^\infty(C_\tau)} &= O(\tau^{-1/2-h}). \end{aligned}$$

The core is defined by fixed similarity coordinates. It is not a material parcel of fixed fluid particles. Its complement includes an exterior and annular corrections; the leading core is not the complete solution.

Choose a smooth cutoff χ_τ , between zero and one, supported in a fixed compact inner similarity-coordinate region and equal to one near $(X_*, \eta) = (X_*, 0)$ with $0 < X_* < X_c$. Define

$$v_\tau := \chi_\tau u^{(0)}(1 - \tau).$$

On this support, volume is $O(\tau^{3/2-h})$ and the velocity bound is $O(\tau^{-1/2-h})$. Direct multiplication therefore gives

$$\|v_\tau\|_2^2 = O(\tau^{1/2-3h}), \quad \|v_\tau\|_2 = O(\tau^\gamma), \quad \gamma := \frac{1}{4} - \frac{3}{2}h > 0.$$

The lower bound is a separate source input: [6, Section 3.1, p. 8, the displayed azimuthal velocity on the circle $z = 0$, $r = \sqrt{2X_*\tau}$] gives $E(X_*, 0) > 0$ and

$$u_\theta^{(0)}(\sqrt{2X_*\tau}, 0, 0, 1 - \tau) = E(X_*, 0)\tau^{-1/2-h}.$$

Here the arguments are (r, θ, z, t) . Since $\chi_\tau = 1$ there, $\|v_\tau\|_\infty \geq E(X_*, 0)\tau^{-1/2-h} \rightarrow \infty$; no lower bound is inferred from the preceding O -estimates. Proposition 4.2 simultaneously yields

$$\sup_{\phi \in \mathbb{P}} |\langle v_\tau, \phi \rangle| = O(B\tau^\gamma) \rightarrow 0$$

for every uniformly L^2 -bounded probe family. In particular,

$$\|S_q v_\tau\|_\infty \leq C q^{-3/2} \tau^\gamma \rightarrow 0 \quad \text{for fixed } q > 0.$$

These are statements about the defined leading-core contribution. They are not a proof that the full flow approaches zero or a smooth background. The complete field is instead covered by Theorem 5.1 and Corollary 5.2 through the full hypothesis check at the start of this section.

6.2 A correction concerning local vorticity

Cutting off velocity and cutting off vorticity are different operations:

$$\nabla \times (\chi_\tau u) = \chi_\tau \omega + \nabla \chi_\tau \times u.$$

A calculation using $\chi_\tau \omega$ cannot silently replace it by $\nabla \times v_\tau$. Nor can a bound on one radial derivative be substituted for a bound on the complete Cartesian gradient without controlling every component. The velocity-based estimate above avoids these gaps.

For the cutoff field itself, Proposition 4.1 remains valid distributionally and gives

$$\|K_Q^*(\nabla \times v_\tau)\|_\infty \leq C Q^{-5/2} \tau^\gamma \rightarrow 0.$$

No claim is made that v_τ separately satisfies incompressibility or the Navier–Stokes equation. Those properties are not needed for this convolution estimate.

6.3 Why this does not decide BCB

The source exponents do not determine the signed quantities at the maximisers of the mollified field. In particular,

$$\xi_\ell \cdot K_\ell^*(S\omega)$$

cannot generally be replaced by a pointwise projection $\hat{\omega} \cdot S\omega$. The full field includes corrections whose derivatives and quadratic products may affect the result. Moreover, two terms having the same asymptotic order does not establish the coefficient inequality required by $0 < \vartheta < 1$.

Accordingly, the earlier suggested verdict “BCB-compatible but distinguishability-inadmissible” is not established. Neither its first clause nor its second follows from the scaling calculation. The present operational results require no such verdict.

7. An explicit bridge into the existing VERSF metric

7.1 What is given and what must still be identified

Let H_{cl} be the native closure-state Hilbert space and let $\Omega_{\max}$ be the projection used in [5]. A fluid-to-closure map would have the form

$$J_B: \text{fluid and environmental configurations} \rightarrow H_{cl}, \quad T_B = \Omega_{\max} J_B.$$

The symbol B fixes the physical domain, calibration, apparatus, environment, observation duration, and resource budget. These cannot be allowed to change silently during a limiting argument.

The retrieved sources establish the native metric and projection structure, but do not identify this particular Navier–Stokes encoding or prove bounds for every allowed record-producing fluid interaction. The following propositions therefore state **bridge hypotheses**, not additional VERSF axioms and not consequences of finite distinguishability alone.

Proposition 7.1. Bounded finite-rank linear bridge

Fix the environmental context. Suppose a particular instantaneous encoding is a bounded linear map

$$T: L_\sigma^2 \rightarrow A, \quad \dim A = d < \infty.$$

Then the complete flow of Theorem 5.1 satisfies

$$D_{op}(Tu(t), Tu_*) \rightarrow 0.$$

For any solenoidal $v(t) \rightarrow 0$ strongly in L^2 ,

$$\|Tv(t)\|_{op} \leq \|T\| \|v(t)\|_2 \rightarrow 0.$$

Proof. Choose an orthonormal basis $e_1, \dots, e_d$ of A . By the Hilbert-space representation of bounded linear functionals there exist $\phi_j \in L^2_\sigma$ such that

$$Tu = \sum_{j=1}^d \langle u, \phi_j \rangle e_j.$$

Weak L^2 convergence gives convergence of all d coefficients. Finitely many convergent coordinates converge in the Hilbert norm. The second assertion is the operator-norm inequality. ▀

For a cutoff component that is not solenoidal, use the bounded extension TP; the Leray projection satisfies $\|Pv\|_2 \leq \|v\|_2$. This handles the component in Section 6 without assigning a solenoidal-only map to a field outside its domain.

If the coefficient probes also satisfy the derivative conditions of Section 5, a uniform rate and finite-variation bound follow from that section. Finite rank alone gives convergence but does not supply those quantitative rates.

Proposition 7.2. A more general observation-stability bridge

Suppose an encoding, possibly nonlinear, obeys

$$D_{\text{op}}(Tu, Tv) \leq L_B d_\phi(u, v)$$

on the relevant configuration class, including u_* , with fixed $L_B < \infty$. Then its terminal distance and total variation are bounded by L_B times the corresponding bounds in Theorem 5.1. A sequence of sampled encoded states separated successively by $\delta_{\min}$ has

$$N \leq 1 + \frac{L_B L_\phi}{\delta_{\min}}.$$

Proof. Apply the assumed Lipschitz inequality to each two-time difference and to each term of a partition, then use Theorem 5.1 and Corollary 5.3. ▀

This concerns a continuous operational encoding, not a proof that a discrete commitment map is continuous. Thresholding, irreversibility, and environmental memory require their own model. Treating every distinct continuum time as a different committed state would reintroduce the assumption the comparison is intended to examine.

Proposition 7.3. A correct finite-dimensional packing bound

Let operational images lie in a Hilbert ball of real dimension d of radius R , and let Σ_δ be a pairwise- δ -separated set of such images, with $\delta > 0$. Then

$$|\Sigma_\delta| \leq \left(1 + \frac{2R}{\delta}\right)^d.$$

Proof. Open balls of radius $\delta/2$ around the separated points are disjoint and are all contained in the ball of radius $R + \delta/2$. Comparing their Euclidean volumes gives

$$|\Sigma_\delta| \omega_d (\delta/2)^d \leq \omega_d (R + \delta/2)^d.$$

Every finite subset obeys this bound, which also excludes an infinite separated set. ▀

For Proposition 7.1 one can take $R = \|T\|U$. This is a conditional connection between finite energy, a specified finite-dimensional physical encoding, and finite distinguishability. It is not a derivation that all physical encodings have finite rank.

The enclosing neighbourhood is important. The volume argument in [5, Theorem 10.1] places disjoint balls around admissible points but does not establish that those balls lie inside the operational image. For an arbitrary admissible set, its volume alone therefore cannot justify that step. Proposition 7.3 uses an enclosing ball and avoids this boundary gap. Likewise, a bounded region in an infinite-dimensional state norm need

not be totally bounded. An infinite orthonormal sequence has bounded norm and fixed nonzero pairwise separation. The finite-dimensional bridge or an independently established compactness property cannot be omitted.

Corpus correction record VERSF-OG-10.1-20260915 — open in [5]. The preceding boundary gap is a substantive correction to [5, Theorem 10.1, pp. 15–16], not solely to the original Navier–Stokes paper [1]. Proposition 7.3 is the replacement bound to port into [5], with an enclosing radius and a clearly specified separation. The volume-only covering estimate in [5, Corollary 9.2] and repetitions of the original packing formula must also be checked. This revision records the debt but does not alter [5]; the separate correction record [16] contains the proposed replacement and closure criteria, also summarised in Appendix C.

8. Countertests against stronger conclusions

8.1 Individually finite resources are not uniformly bounded resources

Proposition 4.2 contains the quantifier

$$\sup_{\phi \in \mathcal{P}} \|\phi\|_2 < \infty.$$

It does not follow from the assertion that each ϕ separately has a finite norm. For a nonzero $v_\tau \rightarrow 0$ in L^2 , the mathematical probes

$$\phi_\tau = \frac{v_\tau}{\|v_\tau\|_2^2}$$

satisfy $\langle v_\tau, \phi_\tau \rangle = 1$, while $\|\phi_\tau\|_2 = 1/\|v_\tau\|_2 \rightarrow \infty$. Each probe at a positive τ has finite norm. The family does not have a uniform budget.

This is not a construction of a physically realisable detector. It is a counterexample to the logical inference that “finite at every stage” supplies the uniform estimate required for the limit. Physical impossibility would need an independent resource, backreaction, or commitment bound excluding the relevant family.

8.2 Linear readouts do not exhaust nonlinear observables

On the periodic box, the solenoidal sequence

$$w_n(x) = (\sin(nx_2), 0, 0)$$

converges weakly to zero and disappears under every fixed smooth linear probe. Nevertheless,

$$w_n \otimes w_n \rightharpoonup \text{diag}(1/2, 0, 0)$$

in the distributional sense. This sequence is a mathematical counterexample, not an asserted Navier–Stokes time evolution. It shows why linear invisibility cannot be promoted to absence of nonlinear momentum flux.

For the stronger hypothesis $u = b + v$ with $\|v\|_2 \rightarrow 0$ and bounded $\|b\|_2$, an instantaneous quadratic estimate is available:

$$\|u \otimes u - b \otimes b\|_1 \leq 2\|b\|_2\|v\|_2 + \|v\|_2^2 \rightarrow 0.$$

This follows by expanding the tensor product and applying Cauchy–Schwarz. It extends the selected-core argument to bounded stress probes. It does not control derivative-sensitive quantities such as $\|\nabla v\|_2^2$, and it does not upgrade weak convergence of the complete flow to strong convergence.

8.3 An endpoint is not a history

A detector may retain a record of earlier activity even when the instantaneous contribution later becomes small. Heat production, integrated stresses, and emitted signals are history-dependent. Propositions 4.2 and 7.1 do not erase those records.

For a bounded linear history probe with $\|\phi(s)\|_2 \leq B$ and a component satisfying $\|v(s)\|_2 \leq C(T-s)^q$, $q > -1$, one can bound only its late-time tail:

$$\left| \int_{T-\varepsilon}^T \langle v(s), \phi(s) \rangle ds \right| \leq \frac{BC}{q+1} \varepsilon^{q+1}.$$

Earlier accumulated contributions remain. A complete VERSF test must therefore specify whether it concerns instantaneous configurations, complete histories, or already committed environmental records.

8.4 A measurement threshold is not an equivalence relation

For a general metric, the relation $d(u, v) < \delta$ need not be transitive. It cannot simply be declared to define quotient classes. Exact equality of all specified readouts does define an equivalence relation; alternatively, one may use an independently defined record partition. Packing arguments need only pairwise separation and do not require a spurious threshold quotient.

8.5 The finite-dimensional metric space is not the discrete fact set

The continuous space A of [5] supports arbitrarily close mathematical points. Its discrete, ledger-admissible subset is a different object. Finite dimension alone is not a claim that every mathematical point is a Fact, nor that every mathematically refined trajectory requires an unbounded sequence of independently recorded Facts. The admissibility and commitment map must do that work.

8.6 Small adjacent refinements do not establish indistinguishability of the tail

Section 7 of the quantitative draft [10] suggested testing successive refinements by requiring their observed differences to fall below a tolerance. That is not a sufficient endpoint criterion. For example,

$$a_n = \log(n+1), \quad |a_{n+1} - a_n| = \log\left(1 + \frac{1}{n+1}\right) \rightarrow 0,$$

although $a_n \rightarrow \infty$. Arbitrarily small steps can accumulate into an arbitrarily large change. Even convergence alone does not guarantee finite variation: $a_n = (-1)^n/n$ tends to zero but has divergent total successive variation.

The replacement is the two-time estimate already proved in Theorem 5.1. Writing

$$G_\phi(s) = B_1 U^2 + \nu B_2 U + B_0 \|f(s)\|_2,$$

we have $G_\phi \in L^1(0, T)$ and, for all $t_0 \leq t < s < T$,

$$d_\phi(u(t), u(s)) \leq \int_t^s G_\phi(a) da \leq \int_{t_0}^T G_\phi(a) da \rightarrow 0 \quad (t_0 \uparrow T).$$

This controls the **entire late-time tail**, not merely adjacent samples. The same integrable bound gives finite total variation and the separated-readout bound in Corollary 5.3. Transferring it to physical records still requires the channel assumptions in Section 7. Model validity is tested separately in Section 10; it is not folded into a new definition of distinguishability.

9. The pulse construction and its quantitative scale hierarchy

9.1 Inverse stress matching is not, by itself, circular reasoning

The quantitative draft [10, Section 3] describes the use of two pulse families to supply a prescribed mean momentum flux. The relevant source is [6, Proposition 7.5, pp. 82–83]. In the notation of that proposition, the covariance of two disjoint pulse families has the form

$$C(a_+ b_+ + a_- b_-) = H \begin{pmatrix} a_+^2 \\ a_-^2 \end{pmatrix}.$$

The target vector is prescribed by the background residual. If the covariance matrix H is invertible and $H^{-1}\mathbf{T}$ has positive entries, then $a_\sigma = \sqrt{(H^{-1}\mathbf{T})_\sigma}$ supplies real amplitudes. Invertibility, positivity, regularity and compatibility with the pulse equations are the substantive proof obligations; adjustable amplitudes alone are not evidence that those obligations have been assumed.

We retain that logical interpretation, not a verdict on the entire proof. In particular, mathematical selection of internal pulse amplitudes is not evidence that an external actuator must individually impose every later pulse. The source's pulses exchange energy with background shear; its smooth force is distinct from the internally amplified motion [6, equation (7.22)]. Neither robustness under perturbation nor lack of robustness follows simply from the inverse-designed form.

9.2 Source inputs and the carrier wavelength

The following are the source statements used for the scale calculation, rather than additional VERSF assumptions. The dyadic chart uses $Q = 2^{-j}$ and $\epsilon = Q^h$, with fixed $0 < h < 1/100$ [6, equation (6.1), p. 62]. We write the band index as j to distinguish it from the original paper's smoothing length ℓ . The carrier integer is $k = \lceil \epsilon^{-1/2} \rceil$ [6, equation (7.2), p. 74]. The evaluated chart derivatives satisfy $\nabla_* = \sqrt{Q} \nabla_x$, and the phase normal $n_\phi = \nabla_* \Phi$ has positive uniform lower and upper bounds in the active pulse rectangles [6, equation (6.6), p. 63; equations (7.4) and (7.9), pp. 74–75].

For the fundamental carrier $e^{ik\Phi}$, define its local physical wavevector and wavelength by

$$\mathbf{k}_p = \nabla_x(k\Phi) = \frac{k}{\sqrt{Q}} n_\phi, \quad \lambda_p = \frac{2\pi}{|\mathbf{k}_p|}.$$

For $0 < Q \leq 1$,

$$Q^{-h/2} \leq k \leq Q^{-h/2} + 1 \leq 2Q^{-h/2}.$$

It follows, with constants independent of the late band, that

$$|\mathbf{k}_p| \asymp Q^{-(1+h)/2}, \quad \lambda_p \asymp Q^a, \quad a := \frac{1+h}{2} > 0.$$

Since the radial scale satisfies $\ell_r \asymp Q^{1/2}$ on the relevant similarity region,

$$\frac{\lambda_p}{\ell_r} \asymp Q^{h/2} \rightarrow 0.$$

This is a local carrier law, not an identification of the complete field's smallest scale, its entire Fourier spectrum, or an actuator bandwidth. Higher harmonics and support envelopes must be tracked separately. The carrier alone already provides a scale tending to zero. The law and its interpretation are the quantitative draft's main useful addition [10, Section 5.1]; the calculation above makes the ceiling function and the source loci explicit.

Proposition 9.1. A positive material validity scale permits only finitely many carrier bands

Fix an active reference band Q_0 and a dimensional reference wavelength $\lambda_{p0} > 0$. Let $Q_n = Q_0 2^{-n}$, $n \geq 0$. Suppose the active carrier wavelengths satisfy uniform bounds

$$c_\lambda \lambda_{p0} 2^{-an} \leq \lambda_p(Q_n) \leq C_\lambda \lambda_{p0} 2^{-an}, \quad 0 < c_\lambda \leq 1 \leq C_\lambda < \infty.$$

If a specified material approximation requires $\lambda_p \geq \delta_{\text{mat}}$ for a fixed $\delta_{\text{mat}} > 0$, every band satisfying that requirement must obey

$$n \leq \frac{1}{a} \log_2 \left(\frac{C_\lambda \lambda_{p0}}{\delta_{\text{mat}}} \right).$$

Define $n_{\text{cap}} = a^{-1} \log_2(C_\lambda \lambda_{p0}/\delta_{\text{mat}})$. The number of such bands, counting the reference band when applicable, is therefore at most

$$N_{\text{bands}} \leq \max(0, 1 + \text{floor}(n_{\text{cap}})).$$

Here $\text{floor}(x)$ denotes the largest integer no greater than x .

Proof. Combine $\delta_{\text{mat}} \leq \lambda_p(Q_n)$ with the upper wavelength bound and solve the resulting inequality for n . All larger integers fail the requirement. No monotonicity of the hidden profile prefactors is needed. ▀

This is a bound on the number of **scale levels** above the stipulated material threshold, not the number of individual pulses, independently encoded Facts, or independent microscopic states. It is not a physical proof that smaller structures disappear. It says that following those structures with the same material approximation would violate its specified validity rule.

9.3 Crossover scales, constants, and illustrative band counts

Put $s = Q/Q_0$. The corresponding continuous comparison is

$$c_\lambda \lambda_{p0} s^a \leq \lambda_p(Q) \leq C_\lambda \lambda_{p0} s^a.$$

A guaranteed failure occurs whenever

$$s < s_\delta^- := \left(\frac{\delta_{\text{mat}}}{C_\lambda \lambda_{p0}} \right)^{1/a}.$$

Conversely, when it lies inside $0 < s \leq 1$, the inequality

$$s \geq s_\delta^+ := \left(\frac{\delta_{\text{mat}}}{c_\lambda \lambda_{p0}} \right)^{1/a}$$

is sufficient for this one wavelength requirement to hold. There may be a transition region between these bounds. The attachment's notation

$$Q_\delta = Q_0 \left(\frac{\delta_{\text{mat}}}{\lambda_{p0}} \right)^{2/(1+h)}$$

is therefore an equality-scale estimate, not an exact calibrated first-crossing time. With uniform constants, the band index differs from $(2/(1+h)) \log_2(\lambda_{p0}/\delta_{\text{mat}})$ by a bounded amount. If the constants or validity scale vary with the band, that variation must be included instead.

For illustration only, take $h = 0.005$ and unit prefactors. Wavelength-to-threshold ratios of 10^6 and 10^{12} give real-valued additional-band indices approximately 39.66 and 79.33. In the exact normalised toy law, bands $n = 0, \dots, 39$ and $n = 0, \dots, 79$, respectively, satisfy the threshold. These are checks of the algebra, not predictions for a selected fluid or the uncalibrated external construction.

9.4 Relating bands to the time remaining

In a fixed inner similarity region bounded away from the chart's edge, the source has $q_{\text{src}} \asymp \tau = 1 - t$, and active dyadic bands satisfy $q_{\text{src}} \asymp Q$ [6, Section 3.1, p. 8; Section 6.2, p. 64]. This relation is local to the region being followed; it is not an assertion that every pulse anywhere in the field has $Q = \tau$.

With a fixed dimensional calibration and reference remaining time τ_0 , the corresponding local law is

$$\frac{\lambda_p(\tau)}{\lambda_{p0}} \asymp \left(\frac{\tau}{\tau_0} \right)^a, \quad \tau_\delta \asymp \tau_0 \left(\frac{\delta_{\text{mat}}}{\lambda_{p0}} \right)^{1/a} > 0.$$

Uniform comparison constants give upper and lower crossover bounds just as above. A physical number of seconds requires a material and a dimensional calibration; the normalised singular time $t = 1$ is not

automatically one second. The result is that a fixed positive threshold is encountered at a positive time remaining, not a universal prediction of when this happens.

10. Material-model validity: conditional stopping conditions

10.1 Three different meanings of resolution

The measurement scale q , the material validity scale δ_{mat} , and VERSF's operational separation δ_{min} have different roles. The first belongs to a chosen averaging procedure. The second specifies when a particular constitutive or coarse-grained material description is reliable. The third belongs to the framework's state-space distinguishability structure. No equality between them is assumed.

For a fixed material approximation, δ_{mat} might encode a molecular averaging requirement, a correlation or relaxation scale, or a noise-based tolerance. Specifying it is a physical input, not a consequence of finite energy or of Proposition 9.1. A numerical grid spacing would instead supply a computational resolution limit; it cannot by itself establish a material limit.

10.2 Dilute-gas Knudsen comparison

For a dilute gas, the Knudsen number compares molecular mean free path with a relevant flow scale. The small-scale kinetic literature studies where Navier–Stokes constitutive descriptions cease to be adequate and kinetic corrections are required [11]. Using the carrier as the selected variation scale, define

$$\text{Kn}_p(Q) = \frac{\lambda_{\text{mfp}}(Q)}{\lambda_p(Q)}.$$

Assume, for the regime being tested, that $\lambda_{\text{mfp}}(Q) \geq \lambda_* > 0$ and that the material model requires $\text{Kn}_p \leq \text{Kn}_{\text{crit}}$. The upper wavelength bound gives

$$\text{Kn}_p(Q) \geq \frac{\lambda_*}{C_\lambda \lambda_{p0}} s^{-a}.$$

Consequently, the selected Knudsen condition must fail when

$$s < s_{\text{Kn}} := \left(\frac{\lambda_*}{C_\lambda \lambda_{p0} \text{Kn}_{\text{crit}}} \right)^{1/a}.$$

When the mean free path can be treated as constant and the prefactors are normalised, this recovers the attachment's estimate

$$Q_{\text{Kn}} \asymp Q_0 \left(\frac{\text{Kn}_{p0}}{\text{Kn}_{\text{crit}}} \right)^{2/(1+h)}.$$

The positive lower bound on mean free path is essential to the guaranteed divergence. If the mean free path changes, the relevant test is its ratio to λ_p , not the word “molecular” by itself. For example, if $\lambda_{\text{mfp}} \asymp s^\eta$, then $\text{Kn}_p \asymp s^{\eta-a}$; divergence follows from this comparison only when $\eta < a$.

No universal numerical value of Kn_{crit} is assigned. Validity depends on the observable and tolerance. McMullen and colleagues [12] compared deterministic fluid calculations with molecular-gas-dynamics simulations of a decaying Taylor–Green vortex and found differences associated with thermal fluctuations at small turbulent scales, including scales larger than the mean free path. That is evidence about their simulated regime, not a verification of this singular construction and not a universal cutoff theorem. We do not describe that simulation study as a laboratory experiment.

Liquids require their own molecular, correlation and relaxation criteria; the dilute-gas mean-free-path formula is not imported unchanged into a liquid calculation.

10.3 Low-Mach validity and the velocity being compared

Low-Mach limits provide a mathematical connection between compressible viscous equations and incompressible flow under specified hypotheses [13]. The Mach number compares a characteristic flow speed with sound speed; density changes are part of the physical issue it diagnoses [14]. We use it as a **stipulated validity condition for a chosen low-Mach approximation**, not as a theorem that every high-speed incompressible solution is physically impossible.

Let $W(Q)$ be a characteristic velocity difference or ambient-relative speed relevant to that approximation. It is not the global L^2 bound U used earlier. A uniform change of inertial frame must not manufacture a model failure: the fluid and reference velocity are transformed together.

The source supplies a nonzero speed on a shrinking core path for the complete constructed field,

$$u_\theta(x_\tau, 1 - \tau) = \tau^{-b} \left(e_0 + O(\tau^{2h}) \right), \quad b := \frac{1}{2} + h, \quad e_0 > 0,$$

[6, equations (10.20)–(10.21), p. 124]. In the source's rest-at-infinity frame this gives the required speed growth. For a material interpretation it can be compared to ambient fluid at rest, making the comparison a relative one. The annular pulses and the shrinking core are different parts of the same construction; this speed estimate is not obtained from the pulse wavelength alone.

After choosing physical units and a reference speed W_0 , assume the corresponding lower bound is

$$W(Q) \geq c_W W_0 s^{-b}, \quad c_W > 0.$$

Suppose also that $c_s(Q) \leq c_{s,\max} < \infty$ throughout the regime being invoked, and that the approximation requires $\text{Ma}(Q) := W(Q)/c_s(Q) \leq \text{Ma}_{\text{crit}}$. Then

$$\text{Ma}(Q) \geq c_W \text{Ma}_0 s^{-b}, \quad \text{Ma}_0 := \frac{W_0}{c_{s,\max}},$$

so that requirement must fail when

$$s < \bar{s}_M := \left(\frac{c_W \text{Ma}_0}{\text{Ma}_{\text{crit}}} \right)^{1/b}.$$

For a constant sound speed and two-sided speed comparison, the equality-scale estimate is

$$Q_M \asymp Q_0 \left(\frac{\text{Ma}_0}{\text{Ma}_{\text{crit}}} \right)^{2/(1+2h)}, \quad \ell_{r,M} \asymp \ell_{r0} \left(\frac{\text{Ma}_0}{\text{Ma}_{\text{crit}}} \right)^{1/(1+2h)}.$$

With only an upper bound on sound speed, $\bar{s}_M$ is a sufficient late-time failure threshold, not an exact estimate of the earliest failure. If $c_s \asymp s^{-\zeta}$ instead, the comparison gives $\text{Ma} \asymp s^{\zeta-b}$ and forces growth only when $\zeta < b$. Selecting a thermodynamic evolution that changes sound speed requires additional physics and cannot be silently incorporated into the original incompressible description.

Proposition 10.1. Exit from the stipulated material regime

Fix a dimensional realisation, an active similarity region and a material model whose validity requires the applicable conditions in Sections 10.1–10.3. Under their respective uniform assumptions, each positive threshold $\bar{s}_\delta$, $\bar{s}_{\text{Kn}}$ or $\bar{s}_M$ supplies a sufficient condition for failure of its corresponding validity requirement. Hence the intersection of those requirements cannot remain satisfied down to $s = 0$.

Proof. Each relevant ratio diverges, or wavelength vanishes, by the displayed bounds. Below the associated positive threshold, at least one required inequality fails. With several applicable criteria, $s < \max\{\bar{s}_\delta, \bar{s}_{\text{Kn}}, \bar{s}_M\}$ ensures that at least one fails. The local relation $Q \asymp \tau$ then gives a positive remaining-time bound. ▀

The maximum is a combined **guaranteed-failure bound**, not an exact first-exit time without additional calibration and monotonicity. A criterion already violated at the reference band provides no later validity

interval. The molecular-scale and Knudsen tests may encode the same physical limitation and must not be counted as independent experimental evidence.

The proposition means that a material cannot both follow the specified scale laws indefinitely and remain inside the **stipulated approximation regime**. It does not determine whether a more complete model smooths the flow, forms a shock, changes phase, or develops another singular structure. Nor does a sub-threshold wavelength alone establish a large error in every coarse observable; amplitudes, tolerances and nonlinear transfers matter.

10.4 Smooth forcing, internal motion, and physical implementation

The external theorem's smooth forcing is not an infinite kick. Smoothness also does not specify the response bandwidth or noise tolerance of a material implementation. Conversely, a mathematically fine correction need not be a separately controlled actuator input: the construction uses internal shear amplification. An implementation claim would require a specified driving mechanism, its coupling to the fluid, tolerances, and a test of the evolution under those tolerances. None is supplied by solving the algebraic stress-matching system alone.

A finite-resolution or finite-band numerical calculation must eventually stop resolving the exact hierarchy. That is a limitation of that calculation, not a numerical disproof of the continuum theorem. A physical multiscale calculation must instead state the material criteria used to replace the continuum model and must evaluate the resulting motion in the replacement model. The present paper supplies crossover formulas and diagnostics, not that simulation.

11. The physical connection to VERSF and a test programme

11.1 Keep the three conclusions distinct

The readout result is analytic: a specified common probe budget produces convergent measurements and finite variation. The material result is conditional: a specified continuum or low-Mach validity rule is eventually violated by the quoted scale hierarchy. The VERSF question concerns which differences are converted into separately stabilised records through the physical fluid–environment interaction. No one of these statements implies the other two.

In particular, model failure is not automatically a violation of finite distinguishability. The molecular or kinetic dynamics that replace the approximation may remain perfectly meaningful. Conversely, a finite readout does not show that no other record was formed. The operational threshold is already part of the VERSF framework; there is no need to invent a new axiom to obtain either result.

11.2 Specify the physical comparison before selecting a norm

A direct test must fix a material, forcing mechanism, dimensional calibration, region, environment, relevant alternatives and observation interval. Comparing a flow with a cutoff component removed is not the same as comparing two actual solutions with the same forcing and initial data. The latter would require its own existence and stability argument.

The existing VERSF metric and projection can be used only after supplying the physical encoding J_B of Section 7. A bounded finite-rank instantaneous map gives Proposition 7.1; a justified observation-stability inequality gives Proposition 7.2. Neither property is established merely by stating that all experiments are finite.

11.3 Track the material limits and the record-producing channels separately

For the material trajectory, track the carrier and envelope scales, molecular or correlation scale, relevant relative speed, sound speed, microscopic population per averaging region, thermal fluctuations and driving tolerances. The cutoff formulas apply only while their comparison constants and material assumptions hold. When a criterion fails, a compressible, kinetic, molecular or fluctuating model must be chosen and checked; it is not legitimate simply to continue calling the same calculation a validated material-fluid prediction.

For record production, track sensitivity, amplification, noise, duration, backreaction and memory, including nonlinear and derivative-sensitive interactions. Theorem 5.1 controls the stated solenoidal linear probes. Section 8 explains why that does not exhaust heat, stress, nonlinear response or already accumulated records. A proof of a uniform physical channel budget must address the actual class, not exclude a difficult channel by choosing a weaker norm.

11.4 What would constitute progress

A physically justified uniform bound would transfer the relevant fluid result into the VERSF operational description. A demonstrated violation of an independently specified VERSF bound would establish exclusion within that framework. A successful multiscale continuation would instead show what the chosen material does after the incompressible approximation fails. These are different achievements.

An experimental test of VERSF requires a quantitative prediction distinguishing its proposed record formation from alternatives. Observing ordinary molecular or compressible corrections alone would not uniquely establish VERSF. This manuscript reports neither such an experiment nor a completed full-field simulation.

12. Conclusions

The original VERSF manuscript asked how a vortex might intensify and whether the continuum's ever-finer distinctions must be physically realisable. Its disputed proof steps remain subject to Appendix A. The present analysis answers more limited, precisely stated questions without relying on those steps.

The analytic result. Finite-energy Navier–Stokes flow satisfying the stated weak identity and integrable-force assumption has a unique weak terminal state. Every probe family obeying the specified common bounds has convergent readouts and finite total variation. Fixed positive-resolution velocity and vector-mollified-vorticity measurements are included. A vanishing-energy concentrating component can be negligible to those probes despite a divergent pointwise amplitude.

The quantitative material result. The cited pulse variables give a carrier law $\lambda_p \asymp Q^{(1+h)/2}$. A specified positive material validity scale permits only finitely many carrier bands. With the additional, explicit mean-free-path or sound-speed bounds, the selected Knudsen or low-Mach regime must also be left before the formal singular limit. These are conditional limits of a material approximation, not a universal minimum length or a prediction of what a replacement theory does.

The remaining physical question. Theorem 5.1 paired with Section 8.1 is the principal bridge: finite variation holds for uniformly bounded probe families, while individually finite probes need not have a family-wide bound. Does the relevant record-forming channel class possess the required uniform budget? The native VERSF metric exists; the fluid-to-record encoding and its physical bounds remain to be established.

Finite distinguishability is not BCB. A material validity threshold is not an operational-state threshold. Small adjacent refinements do not prove that an entire late-time trajectory is indistinguishable. None of these distinctions may be removed merely to obtain a desired verdict.

A correct continuum blow-up theorem, finite behaviour of a specified measurement class, and an eventual exit from a chosen material approximation can all be true at once.

This is the complementary conclusion. The external mathematical claim is judged by its proof and the official conditions; the readout results by their estimates; the material interpretation by its stated regime; and VERSF by the physical connection and discriminating predictions it establishes.

Appendix A. Targeted mathematical audit of the original paper

This appendix concerns the supplied 48-page version of [1]. Page numbers refer to that version. It is a targeted dependency audit, not a claim that every statement in the entire VERSF corpus has been verified. The present paper does not rely on the problematic claims listed here.

A.1 Fixed-scale mollified blow-up: the limit order must be corrected

Page 6 asserts that vorticity blow-up forces $M_\ell \rightarrow \infty$ at all sufficiently small fixed ℓ . For finite-energy flow this is incompatible with Proposition 4.1:

$$M_\ell(t) \leq C_K \ell^{-5/2} \|u(t)\|_2.$$

The relevant unresolved behaviour is in a scale that shrinks with time, or a supremum over scales, not divergence at each fixed positive scale. This correction affects the general physical argument as well as its technical interpretation.

A.2 Lemma 4.1: a different kernel was substituted without an estimate

The transport commutator estimate on page 19 introduces the kernel $|z| |\nabla K(z)|$ and then replaces its convolution by the convolution with K . That replacement requires a domination or comparison estimate which is not provided for the stated general mollifier.

A valid bound is

$$\|[K_\ell^*, u \cdot \nabla] \omega\|_\infty \leq \|\nabla u\|_\infty \|H_\ell^* \omega\|_\infty, \quad H_\ell(z) = |z| |\nabla K_\ell(z)|.$$

Since $\|H_\ell\|_1 = C_1$, a coarser valid estimate is

$$\|[K_\ell^*, u \cdot \nabla] \omega\|_\infty \leq C_1 \|\nabla u\|_\infty \|\omega\|_\infty.$$

Neither estimate, without extra structure, supplies the original closure in M_ℓ and $\tilde{M}_\ell$.

A.3 Lemma 4.2: the norm replacement reverses the necessary inequality

The same page proves a bound involving $\|\omega\|_\infty$ and then substitutes the smaller quantity M_ℓ . The valid consequence is

$$v \|\Delta_x (K_\ell^* \omega)\|_\infty \leq C_2 v \ell^{-2} \|\omega\|_\infty.$$

The fact that $M_\ell \leq \|\omega\|_\infty$ cannot justify replacing the right-hand side by $C_2 v \ell^{-2} M_\ell$. A separate inverse or band-limiting estimate would be needed. Thus Theorem 4.3 is not established by the supplied lemmas as written.

A.4 The moving-scale equation includes an extra term

Theorem 3.1 uses fixed ℓ , while Theorem 5.1 subsequently takes $\ell = \ell(t)$. Let

$$z(t, x) = K_{\ell(t)}^* \omega(t, x), \quad M(t) = \|z(t)\|_\infty.$$

Assume $\ell(t) > 0$ is prescribed and continuously differentiable. At times when $M > 0$ and the maximum-envelope argument applies, choose a maximiser x_t and direction $\xi_t = z(t, x_t)/M(t)$. The exact mollified evolution includes

$$\partial_t z = K_\ell^*(S\omega) - (u \cdot \nabla)z - [K_\ell^*, u \cdot \nabla]\omega + v \Delta_x z + K_\ell^*(\nabla \times f) + \dot{\ell} \partial_\ell K_\ell^* \omega.$$

The projected transport term vanishes at the maximum of $|z|$. Consequently, at almost every differentiability time, the safe lower estimate is

$$\begin{aligned} M'(t) &\geq I_{\text{stretch}} - |[K_\ell^*, u \cdot \nabla]\omega(x_t)| - v |\Delta_x z(x_t)| \\ &\quad + \xi_t \cdot K_\ell^*(\nabla \times f)(x_t) - |\dot{\ell}| |\partial_\ell K_\ell^* \omega(x_t)|. \end{aligned}$$

The scale derivative has the elementary bound

$$\|\partial_\ell K_\ell^* \omega\|_\infty \leq \frac{C_{K,S}}{\ell} \|\omega\|_\infty, \quad C_{K,S} = \|3K + z \cdot \nabla K\|_1.$$

A formula defining ℓ implicitly through M_ℓ also requires existence and regularity of that scale choice. It is not supplied merely by writing $\ell = (v/M_\ell)^{1/2}$.

For smooth forcing, the projected force term satisfies $|I_f| \leq \|\nabla \times f\|_\infty$. Its boundedness is valid, but does not control the other terms or prove BCB.

A.5 Theorem 5.1 does not close its Riccati comparison

On page 21, the proof seeks to choose an initial-amplitude threshold so that

$$C_2 + C_1 C \kappa^2 \leq c/4.$$

The displayed constants are independent of that initial amplitude, so raising the initial amplitude cannot enforce this condition. In addition, the logarithmic error coefficient grows with M . Controlling it only until M reaches a finite multiple of its initial value does not justify comparison all the way to infinity.

A correct conditional comparison can instead be stated as follows: if the **complete** corrected right-hand side, including force and scale-motion terms, is at least aM^2 for a fixed $a > 0$ throughout a sufficiently long interval, then

$$\frac{1}{M(t)} \leq \frac{1}{M(t_0)} - a(t - t_0)$$

as long as the smooth regime and hypotheses persist. Smooth continuation with those hypotheses is impossible beyond $t_0 + [aM(t_0)]^{-1}$. This follows by differentiating $1/M$ almost everywhere and integrating. It does not prove that any actual solution sustains the required net-surplus condition.

Page 31 contains a separate arithmetic error: with $T_R = 2/(c_0 M(0))$, the comparison time $4/(c_0 M(0))$ is greater than T_R , not less. The downstream persistence theorem must therefore be revisited independently of the other corrections.

A.6 Direction coherence needs a quantitative mass condition

Pages 26–27 replace a low-vorticity contribution by an asserted $O(1 - \lambda)$ error. A high-vorticity threshold by itself does not bound the fraction of the convolution mass outside the high-vorticity set.

A correct local version is available. Let the high-set vorticity make angle at most $\delta_0 < \pi/2$ with a reference direction, and suppose the low-set fraction of $K_\ell^*|\omega|$ is at most ε . Projection onto that reference direction yields

$$|K_\ell^* \omega| \geq [\cos \delta_0 - (1 + \cos \delta_0)\varepsilon](K_\ell^*|\omega|).$$

The bracket must be positive. Evaluating this bound at a maximiser of $K_\ell^*|\omega|$ provides a valid comparability condition. It does not follow merely by choosing λ close to one. The endpoint-control templates also need their precise norm hypotheses; bounded mean oscillation alone is not a supplied L^∞ gradient estimate.

A.7 Trigger occurrence and mechanism failure were treated as equivalent

The proof on page 36 establishes, at most, the implication that absence of all stated triggers entails persistence, provided the preceding quantitative trigger estimates are valid. This makes persistence and trigger occurrence exhaustive alternatives; it does not make them mutually exclusive.

A trigger defined by a sufficient diagnostic threshold may occur even while the desired estimates remain valid. The reverse implication “trigger occurs, therefore A or C fails” is not proved by the displayed contrapositive. Moreover, failure of this particular sufficient blow-up test is not a proof of regularity or proof that every singular mechanism is blocked.

A.8 The time-parametrisation conclusion needs separate repair

Theorem 7.2 on pages 23–24 uses an additional condition $V_{\text{eff}} = \|\omega\|_2^2 / \|\omega\|_\infty^2 \geq V_* > 0$ and asserts the blow-up rate $M \sim (T_* - t)^{-1}$. A one-sided inequality $M' \geq aM^2$ does not prove this asymptotic rate. For example, sufficiently near T_* , $M(t) = (T_* - t)^{-1/2}$ satisfies such an inequality while having a different rate.

There is also an immediate energy consistency check. If both $V_{\text{eff}} \geq V_* > 0$ and $\|\omega\|_\infty \geq c/(T_* - t)$ held near T_* , then

$$\int^{T_*} \|\omega(t)\|_2^2 dt \geq c^2 V_* \int^{T_*} (T_* - t)^{-2} dt = \infty.$$

This contradicts the finite viscous dissipation of the unforced finite-energy setting of [1]. The proposed combination therefore cannot simply be used to infer a physical failure of time. Finally, even a proven divergence of $\int \|\partial_t u\|_2 dt$ would be a statement about path length in that chosen norm, not by itself an identification of physical time with that path length. Section 5 shows how a weaker operational path length can remain finite.

A.9 Audit verdict and non-reliance

The fixed-scale maximum-envelope idea and the elementary energy identity remain useful starting points under appropriate smoothness assumptions. The supplied norm closures, adaptive-scale Riccati theorem, and exclusive outcome dichotomy are not established as written. Merely retaining a “conditional” label does not repair a faulty implication inside its proof.

For that reason, all principal results of this companion paper are proved without invoking Theorem 4.3, Theorem 5.1, or the downstream persistence and time-breakdown theorems of [1]. The original paper should be revised before those results are presented as established.

Appendix B. An exact concentration benchmark

A simple analytic field checks the distinction between pointwise growth and fixed-resolution response. It is not a simulation of the external construction and is not asserted to solve Navier–Stokes.

For $\varepsilon > 0$ and $0 < a < 3/2$, define

$$v_\varepsilon(x) = \varepsilon^{-a-1}(-x_2, x_1, 0) \exp\left(-\frac{|x|^2}{2\varepsilon^2}\right).$$

Direct differentiation gives $\operatorname{div} v_\varepsilon = 0$. Maximisation and Gaussian integration give

$$\|v_\varepsilon\|_\infty = e^{-1/2} \varepsilon^{-a}, \quad \|v_\varepsilon\|_2^2 = \pi^{3/2} \varepsilon^{3-2a}.$$

Thus the pointwise norm diverges while the energy tends to zero. For the normalised Gaussian kernel

$$K_q(x) = (2\pi q^2)^{-3/2} \exp\left(-\frac{|x|^2}{2q^2}\right),$$

completion of the square yields the exact convolution

$$S_q v_\varepsilon(x) = \frac{\varepsilon^{4-a}}{(\varepsilon^2 + q^2)^{5/2}} (-x_2, x_1, 0) \exp\left(-\frac{|x|^2}{2(\varepsilon^2 + q^2)}\right).$$

Therefore

$$\|S_q v_\varepsilon\|_\infty = e^{-1/2} \frac{\varepsilon^{4-a}}{(\varepsilon^2 + q^2)^2}.$$

For fixed $q > 0$, this tends to zero. For a matched resolution $q = c\varepsilon$ with fixed $c > 0$, it equals $e^{-1/2}(1 + c^2)^{-2} \varepsilon^{-a}$ and diverges. This is an exact demonstration of the resolution dependence, not an appeal to a dimensional guess.

The filtered axial vorticity at the origin is also explicit:

$$[\nabla \times S_q v_\varepsilon]_3(0) = \frac{2\varepsilon^{4-a}}{(\varepsilon^2 + q^2)^{5/2}}.$$

It vanishes at fixed q and diverges at matched resolution. The field's unfiltered gradient energy scales as ε^{1-2a} , illustrating why derivative-sensitive observables need separate estimates.

The following values use $a = 1$, $q = 1$, and $E_\varepsilon = \frac{1}{2} \|v_\varepsilon\|_2^2$. They are evaluations of the exact formulas, rounded for display.

ε	Raw maximum	Filtered maximum	Ratio	E_ε
1	0.606531	0.151633	0.250000	2.784164
1/2	1.213061	0.048522	0.040000	1.392082
1/4	2.426123	0.008395	0.003460	0.696041
1/8	4.852245	0.001148	0.000237	0.348020
1/16	9.704491	0.000147	0.000015	0.174010

Symbolic differentiation confirms the divergence-free identity and the displayed convolution derivatives. Independent one-dimensional Gaussian quadrature, combined by tensor-product factorisation, agrees with the closed convolution formula to relative error below 2×10^{-12} at $a = 1$, $\varepsilon = 0.37$, $q = 0.8$, and $x = (0.6, -0.4, 0.2)$. These checks verify the analytic mechanism and the direction of the estimates. This is not empirical evidence for VERSF, an independent check of the full blow-up proof, or a proof that a physically unrestricted detector cannot resolve the field.

Appendix C. Consolidation, corrections and verification record

C.1 Edition lineage and source roles

Version 2.0 consolidates *Continuum Blow-Up and Finite Operational Distinguishability*, version 1.1 [15], and the supplied quantitative draft [10]. Theorem 5.1, Corollaries 5.2–5.3 and Propositions 4.1–4.2 and 7.1–7.3 retain their numbering. The presentation uses the same resolution symbol q and spatial Laplacian Δ_x . The leading-core decay exponent is renamed γ to avoid confusion with the external source’s variable q_{src} .

The starting exposition and general-reader summary have been rewritten as a single narrative. The inverse stress-matching discussion and scale laws are integrated into Section 9; material-validity calculations appear in Section 10. The old Section 9 test programme and Section 10 conclusions are expanded into Sections 11–12. The separate explanatory commentary remains a separate, unchanged document.

C.2 Changes to the quantitative attachment

Its proposed adjacent-stage Refinement-Distinguishability Criterion is not adopted as an axiom or theorem. Section 8.6 gives explicit counterexamples and uses the existing two-time and variation bounds instead. The attachment’s material-model question is preserved, with its conditional status made precise.

Carrier wavelength is distinguished from the complete field’s spectrum, from envelope scales and from an actuator input. The finite-band calculation now states explicit comparison constants and counts scale levels rather than Facts. Fixed material length, positive lower mean-free-path bound, and finite upper sound-speed bound are separate hypotheses. A bounded sound-speed comparison gives a sufficient failure bound; an exact crossing estimate needs stronger calibration. Velocity is taken relative to a material reference, not an arbitrarily boosted coordinate frame. No universal metre, second, Mach threshold or Knudsen threshold is assigned.

The molecular-gas study [12] is described as a simulation comparison, not a laboratory experiment. The CMI date is taken directly from its page and signed statement as 11 September 2026. The finite-band examples are labelled normalised algebraic examples, not measurements of the cited construction.

C.3 Open corpus correction: Operational Geometry

Record **VERSF-OG-10.1-20260915 remains open in [5]**. Theorem 10.1’s disjoint-ball argument does not establish that the balls lie inside the operational image. Proposition 7.3 supplies a valid enclosing-ball replacement in real dimension d . The source’s volume-only covering formula, dependent references and summaries need corresponding review. The separate record [16] supplies the proposed correction and closure criteria. This consolidated manuscript neither edits [5] nor marks the issue resolved there.

The original Navier–Stokes audit in Appendix A likewise remains open in [1]. Stating the errors here does not repair the source or certify the rest of either source manuscript.

C.4 Checks performed and their limits

The retained exact Gaussian benchmark is supplied with its symbolic and numerical check script. The consolidation adds symbolic exponent checks, inversion checks for wavelength and Mach crossover formulas, direct finite-band count checks for the normalised examples, and the two adjacent-increment counterexamples. These checks concern the displayed identities and logical implications. They are not a direct simulation of the 166-page construction, a proof-assistant verification, or empirical evidence for VERSF.

The source theorem, nonzero core profile, complete-field speed lower bound, dyadic chart scales, phase derivative normalisation, carrier integer, and covariance proposition were checked at the cited locations. The entire external proof was not independently re-proved. Public-source statements are recorded as checked on 15 September 2026. The earlier source-snapshot convention is retained for internal VERSF references without an authorial publication date.

References and source record

Edition convention. RS1 denotes a fixed reference snapshot dated 15 September 2026, assigned for this companion where the supplied manuscript has no explicit authorial version and publication date. It is not an inferred publication date or a claim that the source was revised. Stable identifiers and SHA-256 fingerprints are listed in the accompanying source register; internal upload filenames are not citation identifiers.

[1] Taylor, K. *A Rigorous Reduction of Potential Finite-Time Breakdown in 3D Navier–Stokes*. VERSF manuscript, 48 pp. **VERSF-NS-BASE, RS1, 15 September 2026**. Principal audit locations: pp. 6, 18–24, 26–31 and 36–41.

[2] Taylor, K. *On the Assumption of Distinguishability in Physical Theory: A Constraint-First Perspective and the VERSF Programme*. VERSF manuscript, 27 pp. **VERSF-DIST-ASSUMPTION, RS1, 15 September 2026**. Sections 8.2 and 9.7, pp. 7 and 15–16.

[3] Taylor, K. *Facts as Structural Necessities: On the Co-Dependence of Physical Time and Physical Law on Irreversible Fact Formation*. VERSF manuscript. **VERSF-FACTS-STRUCTURAL, RS1, 15 September 2026**. Section 2, Definitions 2–4; Section 6. Its stated distinguishability structure is used without independently certifying the foundational necessity argument.

[4] Taylor, K. *Proto-Factual Structure and Fact Formation: Interpreting Sub-Coherence-Scale Observations*. VERSF manuscript. **VERSF-PROTO-FACTUAL, RS1, 15 September 2026**. Sections 2–3. The source’s unresolved cross-references are not relied upon; only its explicit statements on proto-factual processes and coherence-scale interpretation are used.

[5] Taylor, K. *Operational Geometry in the VERSF Framework: Entropy-Induced Metrics, Admissible Measure, and the Geometry of Finite Distinguishability*. VERSF manuscript, 20 pp. **VERSF-OP-GEOMETRY, RS1, 15 September 2026**. Definitions 3.4–3.7; Theorems 4.1, 5.1 and 9.1. For the correction required in Theorem 10.1, see Proposition 7.3 and open record VERSF-OG-10.1-20260915. No fluid-state encoding is inferred from its native metric.

[6] OpenAI. *Finite Time Blowup for Navier–Stokes*. 2026, 166 pp. Theorem 1.1, p. 1; Sections 2.1 and 3.1; the positive azimuthal profile evaluation in Section 3.1, p. 8; Lemma 10.4. Accessed 15 September 2026. Source: [OpenAI research manuscript](#).

[7] OpenAI. *On the Navier–Stokes Millennium Prize Problem*. Research announcement, 8 September 2026, updated 10 September 2026. Accessed 15 September 2026. Source: [Official announcement](#).

[8] Clay Mathematics Institute. *Navier–Stokes Announcement*. **11 September 2026**, as displayed in both the page’s date field and the signed statement. Date verified directly on 15 September 2026. Source: [CMI statement](#).

[9] Fefferman, C. L. *Existence and Smoothness of the Navier–Stokes Equation*. Official Clay Millennium Prize problem statement, 6 pp., particularly alternatives C and D. Source: [CMI problem statement](#).

[10] Taylor, K. *Continuum Blow-Up and Physical Admissibility: A Quantitative Refinement-Distinguishability Analysis of the 2026 Forced Navier–Stokes Construction*. Draft dated 15 September 2026, 13 pp. Reference snapshot VERSF-NS-PHYSICAL-QUANT, RS1. Sections 3, 5.1 and 8.2–8.4 are integrated here with revisions; its Section 7 criterion is not adopted.

[11] Hadjiconstantinou, N. G. “The Limits of Navier–Stokes Theory and Kinetic Extensions for Describing Small-Scale Gaseous Hydrodynamics.” *Physics of Fluids* **18**, 111301 (2006). DOI: 10.1063/1.2393436. See also the author’s MIT Nanoscale Gas Transport research page, subsection “Second-order slip,” for the mean-free-path and Knudsen interpretation. This is background for a conditional material criterion, not evidence about the particular singular construction.

[12] McMullen, R. M., Krygier, M. C., Torczynski, J. R., and Gallis, M. A. “Navier–Stokes Equations Do Not Describe the Smallest Scales of Turbulence in Gases.” *Physical Review Letters* **128**, 114501 (2022). DOI: 10.1103/PhysRevLett.128.114501. The reported comparison uses molecular-gas-dynamics simulations of decaying Taylor–Green turbulence.

[13] Danchin, R. “Low Mach Number Limit for Viscous Compressible Flows.” *ESAIM: Mathematical Modelling and Numerical Analysis* **39**(3), 459–475 (2005). DOI: 10.1051/m2an:2005019. Used for the low-Mach limiting framework and its qualified scope, not as a theorem excluding all high-Mach incompressible motions.

[14] NASA Glenn Research Center. *Role of the Mach Number*. Beginners Guide to Aeronautics. Official explanatory page, accessed 15 September 2026. Used for the speed-to-sound comparison and compressibility context.

[15] Taylor, K. *Continuum Blow-Up and Finite Operational Distinguishability: A Companion to the VERSF Navier–Stokes Programme*. Working manuscript 1.1, 15 September 2026, 18 pp. VERSF-NS-COMPANION-1. Source edition for the preserved analytic core and audit.

[16] Taylor, K. *Operational Geometry Correction Record*. Record VERSF-OG-10.1-20260915, 15 September 2026. Open source correction to Theorem 10.1 and dependent volume-based bounds in [5].

Verification boundary. This consolidated draft contains self-contained proofs under stated assumptions and conditional physical comparisons. It does not certify the complete VERSF corpus, independently verify OpenAI’s full construction, establish a material realisation, or report a new physical experiment. Source-paper corrections remain open unless explicitly recorded as resolved in those papers.