

NAVIER–STOKES REGULARITY VIA CRITICAL NORM TRACKING

Anonymous authors

Paper under double-blind review

ABSTRACT

We study 3D incompressible Navier–Stokes flow on a periodic box and design a diagnostic suite that links classical Prodi–Serrin mixed norms, weak- L^p (Lorentz) proxies, scaling-invariant weighted criteria, and sparseness-based geometry. The goal is to evaluate how these regularity indicators behave in direct numerical simulation (DNS) across Reynolds-number sweeps and to quantify numerical sensitivities due to de-aliasing. The study matters for both mathematical regularity theory and practical turbulence modeling because it formalizes computable indicators that can be tracked in simulations and compared to validated datasets. We provide a formal problem setting, an algorithmic pipeline, and a hypothesis-driven evaluation protocol, with explicit acceptance criteria and uncertainty procedures. Empirical results are not yet available because the required scientific Python stack and external dataset access are currently unavailable in the execution environment; we therefore report the full experimental design and the evidence plan that will be used once runs are enabled.

1 INTRODUCTION

The Navier–Stokes regularity problem connects fundamental analysis with predictive modeling of turbulent flows. In practice, DNS and validated databases offer a route to empirically probe critical-norm growth, but bridging the theory–diagnostic gap is nontrivial. Classical Prodi–Serrin mixed-norm criteria provide a baseline for regularity, while weak- L^p and log-improved variants suggest that more permissive integrability can still imply smoothness (Tran & Yu, 2016; Bosia et al., 2014; Bjorland & Vasseur, 2009). Parallel geometric approaches quantify sparseness of intense structures and point to asymptotic criticality, supported by numerical evidence in isotropic turbulence (Grujic & Xu, 2024; Rafner et al., 2021). The present work integrates these lines into a coherent computational protocol that can be run on a periodic box and compared with benchmark data from the Johns Hopkins turbulence database (JHTDB) (Li et al., 2008; Parashar et al., 2024).

Beyond mathematical interest, diagnostics that track near-critical growth have cross-domain relevance: they inform numerical reliability in spectral simulations, support uncertainty-aware turbulence modeling, and provide interpretable indicators for extreme-event forecasting. We focus on the 2π -periodic setting with $2/3$ de-aliasing and a constrained compute budget, reflecting realistic limitations of small-scale experimental studies and open-source pipelines (Patterson & Orszag, 1971; Orszag, 1971).

Contributions.

- We formalize a diagnostic suite that unifies classical mixed norms, weak- L^p (Lorentz) proxies, scaling-invariant weighted criteria, and sparseness metrics in a single evaluation plan.
- We specify a hypothesis-driven evaluation protocol with Reynolds-number sweeps, de-aliasing variants, and explicit acceptance criteria for each diagnostic, enabling reproducible falsification.
- We introduce an aliasing error budget analysis to quantify how high-wavenumber filtering affects near-critical norm spikes.
- We provide a reproducibility plan with seeds, sweeps, compute limits, and uncertainty procedures tailored to periodic DNS and benchmark validation.

2 BACKGROUND AND RELATED WORK

Classical Prodi–Serrin criteria guarantee regularity when the velocity lies in $L_t^p L_x^q$ with $2/p + 3/q \leq 1$ and $q > 3$, and they remain a standard baseline for mixed-norm tracking (Tran & Yu, 2016). Weak- L^p extensions replace

spatial L^q norms with Lorentz $L^{r,\infty}$ norms, widening the admissible class while retaining regularity guarantees (Bosia et al., 2014). Log-improved criteria further relax temporal integrability, suggesting that classical norms may be overly conservative when used as numerical diagnostics (Bjorland & Vasseur, 2009).

Geometric regularity approaches shift attention from integrability to the sparseness of super-level sets of vorticity or higher derivatives, yielding a framework in which the scaling gap closes asymptotically with derivative order (Grujic & Xu, 2024). Numerical evidence in Kida vortex bursts and isotropic turbulence indicates that sparseness exponents can exceed critical thresholds in practice (Rafner et al., 2021). However, comparisons between sparseness indicators and mixed-norm diagnostics are limited, and their relative stability under numerical discretization is underexplored.

On the computational side, pseudo-spectral DNS with 2/3 de-aliasing is a standard technique to control aliasing errors, and additional high-wavenumber filtering has been proposed to further suppress numerical contamination (Patterson & Orszag, 1971; Orszag, 1971). The JHTDB infrastructure supplies validated isotropic turbulence data and web-service access to derivatives and spectra, making it a natural benchmark for diagnostic validation (Li et al., 2008; Parashar et al., 2024). The gap motivating this work is a unified, DNS-computable diagnostic suite that connects these mathematical criteria with realistic numerical constraints and explicit evidence thresholds.

3 PROBLEM SETTING

Let $\mathbb{T}^3 = [0, 2\pi]^3$ be the periodic domain. We study the incompressible Navier–Stokes equations

$$\partial_t u + (u \cdot \nabla)u = -\nabla p + \nu \Delta u, \quad \nabla \cdot u = 0, \quad (1)$$

with initial data $u_0 \in H^1(\mathbb{T}^3)$, viscosity $\nu > 0$, and Reynolds number $\text{Re} = UL/\nu$ with characteristic velocity U and length $L = 2\pi$. The Leray–Hopf energy inequality is

$$\|u(t)\|_2^2 + 2\nu \int_0^t \|\nabla u(\tau)\|_2^2 d\tau \leq \|u_0\|_2^2, \quad (2)$$

which anchors our diagnostic normalization and burst detection (Tran & Yu, 2016).

For Prodi–Serrin pairs (p, q) satisfying $2/p + 3/q = 1$ and $q > 3$, we define the mixed norm

$$\|u\|_{L_t^p L_x^q} = \left(\int_0^T \left(\int_{\mathbb{T}^3} |u(x, t)|^q dx \right)^{p/q} dt \right)^{1/p}. \quad (3)$$

Weak- L^p diagnostics replace the spatial L^q norm with a Lorentz norm

$$\|u\|_{L_t^p L_x^{r,\infty}} = \left(\int_0^T \left(\sup_{\alpha > 0} \alpha \mu(\{x : |u(x, t)| > \alpha\})^{1/r} \right)^s dt \right)^{1/s}, \quad (4)$$

where $3/r + 2/s = 1$ and μ denotes Lebesgue measure (Bosia et al., 2014; Bjorland & Vasseur, 2009). A scaling-invariant weighted diagnostic is defined by

$$W(T) = \int_0^T \frac{\|u(\cdot, t)\|_{L^q}^p}{(1 + S(u(t)))^\gamma} dt, \quad (5)$$

where $S(u)$ is a scaling-invariant norm such as $\|u\|_{L^3}$ or $\|u\|_{\dot{H}^{1/2}}$ and $\gamma > 0$ (Tran & Yu, 2016). For sparseness, define the super-level set of derivatives $D^k u$ by

$$S_k(\lambda, t) = \{x \in \mathbb{T}^3 : |D^k u(x, t)| > \lambda\}, \quad \mu(S_k(\lambda, t)) \propto \lambda^{-\alpha_k}, \quad (6)$$

and compare the estimated exponent α_k to the critical scaling α_k^* implied by the Navier–Stokes scaling $u_\lambda(x, t) = \lambda u(\lambda x, \lambda^2 t)$ (Grujic & Xu, 2024; Rafner et al., 2021).

4 METHODS

4.1 NUMERICAL SETUP AND VALIDATION CONTEXT

We target pseudo-spectral DNS on \mathbb{T}^3 with 2/3 de-aliasing and optional high-wavenumber filtering. This choice is grounded in classical aliasing-control analysis (Patterson & Orszag, 1971; Orszag, 1971) and enables direct comparisons to JHTDB isotropic turbulence spectra and derivative statistics when accessible (Li et al., 2008; Parashar et al., 2024). Reynolds numbers are swept over $\{200, 400, 800, 1600\}$, with resolutions chosen to respect the compute budget while resolving intermittency at a coarse-grained level.

4.2 DIAGNOSTIC COMPONENTS AND MOTIVATION

Classical mixed norms. Prodi–Serrin norms provide the baseline critical-line diagnostic and are the reference for all comparative metrics (Tran & Yu, 2016).

Weak- L^p (Lorentz) proxy. Weak- L^p criteria are theoretically more permissive than Lebesgue-based criteria, motivating a proxy that estimates Lorentz norms from empirical order statistics of $|u|$ (Bosia et al., 2014; Bjorland & Vasseur, 2009).

Scaling-invariant weighting. Weighted diagnostics based on scaling-invariant norms are expected to damp transient spikes while preserving sensitivity to sustained growth (Tran & Yu, 2016).

Sparseness metrics. Super-level set geometry connects directly to asymptotic criticality and provides a complementary, geometric signal for burst dynamics (Grujic & Xu, 2024; Rafner et al., 2021).

Aliasing error budget. Comparing baseline 2/3 de-aliasing to filtered variants quantifies numerical contamination of near-critical spikes, grounding the diagnostic interpretation in numerical analysis (Patterson & Orszag, 1971; Orszag, 1971).

4.3 PIPELINE OVERVIEW

The pipeline consists of (i) DNS generation with controlled de-aliasing, (ii) diagnostic extraction for mixed norms, weak- L^p proxies, and sparseness metrics, (iii) aggregation across Reynolds numbers and seeds with uncertainty estimation, and (iv) validation against benchmark spectra or gradient statistics when dataset access is available (Li et al., 2008; Parashar et al., 2024). Each module is designed to be replaceable: alternative DNS solvers can feed the same diagnostic layer, and validation can use either JHTDB or synthetic proxies when external access is blocked.

Algorithm 1 Critical-norm diagnostic pipeline for periodic DNS.

```

Select  $(p, q)$  pairs on the Prodi–Serrin line and set sweep parameters.
for each Re, resolution, and de-aliasing variant do
  Run pseudo-spectral DNS to produce snapshots of  $u(x, t)$ .
  Compute mixed norms, weak- $L^p$  proxy statistics, and sparseness exponents.
  Detect burst windows via enstrophy or spectral deviations.
  Aggregate metrics across seeds and compute uncertainty estimates.
end for
Compare metrics to acceptance criteria in Table 1.

```

5 EVALUATION PROTOCOL AND PLANNED EVIDENCE

Table 1 summarizes the four hypotheses and the evidence required to validate each claim. The metrics are chosen to directly test whether weak- L^p proxies diverge from classical norms (H1), whether sparseness exponents exceed criticality (H2), whether weighting stabilizes diagnostics (H3), and whether aliasing control bounds spurious spikes (H4). By defining quantitative acceptance criteria, the table serves as the evidence specification for each hypothesis and will be populated once DNS runs are available.

Table 2 records the sweep parameters, seeds, and datasets that will be used to populate the evidence in Table 1. Together, these tables define the experimental design and establish how each hypothesis will be tested and compared across Reynolds numbers and numerical variants.

6 RESULTS STATUS

At present, the results section is restricted to the evidence plan described in Tables 1 and 2. Full DNS runs, figures, and quantitative tables have not yet been generated because the scientific Python stack and external data access required for the diagnostic pipeline are unavailable in the current execution environment. As a consequence, no empirical confirmation or refutation of H1–H4 is reported here. The planned analysis will populate the metrics in Table 1 and will compare the protocol settings of Table 2 against the acceptance criteria once execution is enabled.

Table 1: Hypotheses, diagnostics, and acceptance criteria defining the evidence required to validate each claim. Meeting the criteria provides support for the corresponding hypothesis, while violations indicate falsification or sensitivity to numerical settings.

Hypothesis	Diagnostic evidence	Acceptance criteria (planned)
H1 Weak- L^p divergence	Time series of $\ u\ _{L_t^p L_x^q}$ vs Lorentz proxy; peak-to-mean ratios; cross-correlation and burst alignment.	Lorentz proxy peak-to-mean ratio $\leq 0.9 \times$ classical in $\geq 3/4$ Reynolds numbers and correlation drops below 0.9 in burst windows for ≥ 2 seeds.
H2 Sparseness criticality	Log-log fits of $\mu(S_k(\lambda, t))$ vs λ for $k = 1, 2$; α_k vs α_k^* ; alignment with enstrophy peaks.	Fit $R^2 \geq 0.9$ in $\geq 70\%$ of windows and $\alpha_k > \alpha_k^*$ in $\geq 3/4$ Reynolds numbers for $k = 1, 2$.
H3 Weighted stability	Weighted diagnostic $W(t)$ vs unweighted mixed norm; variance, lead time, and false positive rate.	Variance $\leq 0.85 \times$ unweighted in $\geq 3/4$ Reynolds numbers and lead time improves $\geq 10\%$ without $> 5\%$ FPR increase.
H4 Aliasing error budget	Energy spectra and critical norms under baseline and filtered de-aliasing; spike reduction metrics.	Filtering reduces spike amplitude $\geq 15\%$ with energy decay change $\leq 5\%$ in $\geq 3/4$ Reynolds numbers.

Table 2: Simulation protocol for periodic DNS, including Reynolds numbers, resolutions, seeds, and diagnostic parameter sweeps. These settings specify the experimental conditions used to populate the evidence criteria in Table 1.

Component	Specification
Domain and solver	\mathbb{T}^3 with 2/3 de-aliasing; optional high- k filter variants; pseudo-spectral DNS (Patterson & Orszag, 1971; Orszag, 1971).
Reynolds sweep	$\text{Re} \in \{200, 400, 800, 1600\}$.
Resolutions	$64^3, 96^3$, and 128^3 grids (as budget permits).
Seeds	$\{11, 23, 37, 53\}$ with fixed initialization protocol.
Mixed-norm pairs	$(p, q) \in \{(4, 6), (8, 4), (10, 10/3)\}$ on the critical line.
Weak- L^p proxy	Lorentz exponents $(r, s) \in \{(4, 4), (6, 6)\}$ and sampling stride $\{1, 2, 4\}$.
Sparseness fits	Threshold grids on logspace ranges; derivative orders $k \in \{1, 2\}$; fit windows $\{0.2, 0.3\}$.
Validation datasets	Synthetic DNS and JHTDB isotropic turbulence when accessible (Li et al., 2008; Parashar et al., 2024).

7 LIMITATIONS AND FUTURE WORK

The principal limitation is a data and tooling gap: DNS-derived diagnostics and JHTDB validations cannot currently be executed, and only the experimental protocol is available. This gap limits conclusions to methodological design and precludes empirical validation of the hypotheses in Table 1. It also prevents quantifying uncertainty intervals and numerical sensitivity in a data-driven manner.

7.1 FUTURE WORK

Once the required dependencies and dataset access are available, the immediate follow-up is to run the full Reynolds-number sweep, populate Tables 1 and 2 with measured metrics, and produce the corresponding figures. Additional experiments will refine sampling strides, sparseness fit windows, and filter strengths to assess robustness, and uncertainty will be quantified using repeated seeds and bootstrap confidence intervals.

8 CONCLUSION

We present a unified diagnostic framework for critical norm tracking in periodic DNS that integrates classical mixed norms, weak- L^p proxies, scaling-invariant weighting, and sparseness geometry. The evaluation protocol and acceptance criteria in Tables 1 and 2 define a reproducible evidence plan that will test whether these diagnostics diverge in turbulent bursts and whether numerical aliasing contaminates near-critical spikes. While empirical results are pending, the framework establishes a clear path for bridging theoretical regularity criteria with operational turbulence diagnostics.

REFERENCES

- Clayton Bjorland and Alexis Vasseur. Weak in space, log in time improvement of the ladyzenskaja–prodi–serrin criteria. arXiv preprint, 2009. URL <https://ar5iv.labs.arxiv.org/html/0912.2969>.
- Stefano Bosia, Vittorino Pata, and James C. Robinson. A weak-lp prodi–serrin type regularity criterion for the navier–stokes equations. *Journal of Mathematical Fluid Mechanics*, 2014. URL https://re.public.polimi.it/bitstream/11311/862580/1/A%20Weak-Lp%20Prodi%E2%80%93Serrin%20Type%20Regularity%20Criterion%20for%20the%20Navier%E2%80%93Stokes%20Equations_11311-862580_Pata.pdf.
- Zoran Grujic and Liaosha Xu. Asymptotic criticality of the navier–stokes regularity problem. *Journal of Mathematical Fluid Mechanics*, 2024. URL <https://link.springer.com/content/pdf/10.1007/s00021-024-00888-x.pdf>.
- Yi Li, Eric Perlman, Minping Wan, Yunke Yang, Charles Meneveau, Randal Burns, Shiyi Chen, Alexander Szalay, and Gregory Eyink. A public turbulence database cluster and applications to study lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence*, 2008. URL <https://ar5iv.labs.arxiv.org/html/0804.1703>.
- Steven A. Orszag. On the elimination of aliasing in finite-difference schemes by filtering high-wavenumber components. *Journal of the Atmospheric Sciences*, 1971. doi: 10.1175/1520-0469(1971)028<1074:OTEOAI>2.0.CO;2. URL [https://doi.org/10.1175/1520-0469\(1971\)028<1074:OTEOAI>2.0.CO;2](https://doi.org/10.1175/1520-0469(1971)028<1074:OTEOAI>2.0.CO;2).
- Nishant Parashar, Sawan S. Sinha, and Balaji Srinivasan. Johns hopkins university turbulence database (jhtdb) dataset record. Dataset, 2024. URL <https://doi.org/10.57702/efm5g5uy>.
- G. S. Patterson and Steven A. Orszag. Spectral calculations of isotropic turbulence: Efficient removal of aliasing interactions. *Physics of Fluids*, 1971. doi: 10.1063/1.1693365. URL <https://doi.org/10.1063/1.1693365>.
- Janet Rafner, Zoran Grujic, Christian Bach, Jakob Andreas Baerentzen, Bo Gervang, Ruo Jia, Scott Leinweber, Marek Misztal, and Jacob Sherson. Geometry of turbulent dissipation and the navier–stokes regularity problem. *Scientific Reports*, 2021. URL <https://www.nature.com/articles/s41598-021-87774-y.pdf>.
- Chuong V. Tran and Xinwei Yu. Note on prodi–serrin–ladyzhenskaya type regularity criteria for the navier–stokes equations. Preprint, 2016. URL https://research-repository.st-andrews.ac.uk/bitstream/10023/10048/1/Tran_2016_Criteria_JMP_AAM.pdf.

A REPRODUCIBILITY AND IMPLEMENTATION DETAILS

All runs are planned on a single machine with a total compute budget of at most four CPU hours. The experimental sweep uses Reynolds numbers $\{200, 400, 800, 1600\}$, grid resolutions 64^3 to 128^3 , and seeds $\{11, 23, 37, 53\}$. De-aliasing follows the $2/3$ rule with optional exponential or sharp high- k filtering variants. Diagnostics include mixed-norm pairs $(p, q) \in \{(4, 6), (8, 4), (10, 10/3)\}$, weak- L^p proxy parameters $(r, s) \in \{(4, 4), (6, 6)\}$, and sparseness thresholds on logspace grids with fit-window fractions in $\{0.2, 0.3\}$. Burst detection uses enstrophy growth or spectral deviation thresholds tuned per seed to reduce false positives.

Uncertainty will be estimated via repeated-seed aggregation with nonparametric bootstrap confidence intervals for peak-to-mean ratios, correlation metrics, and sparseness exponents. Approximations include short time horizons and coarse grids to respect the budget; these approximations will be reported alongside all metrics to contextualize sensitivity. Validation against JHTDB will use the published isotropic turbulence dataset and its DOI record when access is enabled (Li et al., 2008; Parashar et al., 2024).