

POTENTIAL MISTRANSLATIONS OF OPENAI'S NAVIER-STOKES NATURAL LANGUAGE PROOF TO LEAN

A. BASTOUNIS, F. CIRCELLI, AND A. C. HANSEN*

1. POTENTIAL MISTRANSLATIONS OF THE NAVIER–STOKES NL PROOF

1.1. AI generated list (with some human input) of suggested mistranslations. This report details an experiment done to compare OpenAI's NL (natural language) proof *Finite time blowup for Navier–Stokes* with the accompanying Lean formalisation at commit `f9e8bc5`. The motivation for this experiment comes from the ideas discussed in the paper *Navier–Stokes lost in translation – Why Lean verification of AI autoformalisation does not guarantee correct natural language proofs*. The twelve comparisons below are generated as a collaboration between ChatGPT and Claude, with some human input.

Warning – Potential hallucinations: Since the list below is generated with the help of AI, and has not been fully manually checked, hallucinations may occur. Hence, the list should be taken as a suggestion of mistranslations that could be checked by others. The purpose of the list is to give an indication of potential mistranslations.

*Corresponding author: a.hansen@damtp.cam.ac.uk.

1. Localization depends on the chosen potential

Paper: Theorem 3.1 represents the local velocity as $u_{\text{loc}} = \text{curl } A + D$, where A is a vector potential and $D = Be_\theta$ is a direct angular field. Equation (10.4) uses this pair to define $u_c = \text{curl}(cA) + cD$, with a smooth cutoff c .

Lean: The local-theorem pair is $A' = A_{\text{raw}} - J$, $D' = D_{\text{raw}} + \text{curl } J$, where J is the axial part removed from the base potential to make A' vanish in the exterior. The compact candidate instead cuts $A_{\text{raw}}, D_{\text{raw}}$. The theorem `velocity_eq_original` identifies the velocities before cutting.

Difference: For the same cutoff, the product rule gives

$$[\text{curl}(cA_{\text{raw}}) + cD_{\text{raw}}] - [\text{curl}(cA') + cD'] = \nabla c \times J.$$

This term does not vanish identically for the selected fields. At $z = 0$, the radial cutoff changes from one to zero for $s = r^2/2 \in (1/64, 1/32)$. At sufficiently late times, the entire segment from s to 1 lies in the heat exterior. There the axial coefficient of J is $-\int_1^s F_{\text{heat}}(v, t) dv$, which is positive because $F_{\text{heat}} > 0$ and $s < 1$. The radial gradient of the cutoff is nonzero somewhere on the transition and is perpendicular to J . Thus the selected localized velocities differ. Their agreement where $c = 1$ and the required properties of the compact candidate are separately proved.

Sources. Theorem 3.1, p. 15; (10.4), p. 118. `LocalPotentialRebundle.lean:24–30`,
`LocalPotentialRebundle.lean:96–109`; `LocalAngularGrowth.lean:138–147`;
`R3CompactCandidate.lean:201–202`; `TailGaugePotential.lean:180–195`, `TailGaugePotential.lean:228–323`;
`SpatialLocalization.lean:32–50`; `RadialHeatProfile.lean:121–122`;
`PhysicalHeatCoordinates.lean:39–44`; `PaperLocalization.lean:28–50`.

2. The cutoff violates the specified domain margin

Paper: The local fields are supplied on $\Omega_* = \{t < 1, q < q_*\}$. Proposition 10.1 chooses the spatial and temporal cutoffs so that $c \neq 0$ implies $q < q_*/2$. This leaves room inside the domain for a smooth extension by zero.

Lean: The cutoffs are fixed independently of q_* . Both equal one at $(t, x) = (3/4, 0)$. The displayed local threshold is

$$q_* = \min \left\{ Q_b, \frac{1}{2(|a_0| + 1)} \right\} \leq \frac{1}{2},$$

where $Q_b > 0$ is a selected band scale and a_0 is the first entry of the cutoff schedule. At that point $q = 1 - t = 1/4$, so $q < q_*/2$ fails.

Difference: This is a direct counterexample to the paper's cutoff-selection condition for the simultaneously displayed Lean choices. The formal construction uses raw fields smooth on all of $\mathbb{R}^3 \times \{t < 1\}$, allowing localization without that margin. The discrepancy is in the domain and extension argument; it does not exhibit a nonsmooth localized field.

Sources. §3.5, p. 16; Proposition 10.1, p. 118. `LocalPaperDomain.lean:16–17`; `SmoothCutoffs.lean:245–265`;
`SpatialLocalization.lean:133–151`; `LocalScaleGeometry.lean:19–23`;
`LocalPotentialRebundle.lean:78–94`.

3. A smooth axis value is replaced by totaled division

Paper: The auxiliary angular-stress quantity Q_s is defined for $X > 0$ by

$$Q_s(X, \eta) = \frac{\int_0^X H(s, \eta) S_q(s, \eta) ds}{XH(X, \eta)}, \quad Q_s(0, \eta) = \frac{1}{2} S_q(0, \eta).$$

Here $H = \sqrt{2X}E = 2XF$, E is the angular velocity profile, F is smooth and nonzero at the axis, and S_q is the smooth source in the radial equation (4.9). The second formula is the smooth extension required after (4.10).

Lean: `ProfileHistories.Profiles.angularLag` uses the quotient literally at every X . At $X = 0$, both numerator and denominator vanish, and Lean's division convention gives 0. The smooth source here is `NaturalEntrance.Sq`; the identity `naturalHistories.angularSource` identifies the integrand in `angularLag` with HS_q . For the selected entrance profile, `source_lower` gives $S_q(0, \eta) > 5/2$ for $|\eta| \leq 1$, so the smooth right-hand limiting value exceeds $5/4$.

Difference: For these entrance data the unweighted functions disagree at the axis, and the raw Lean quotient is discontinuous there. This does not imply discontinuity of the physical stress: Q_s enters it with a vanishing radial factor, and the original code proves the corresponding regular stress identities, including the axis. Regularity of those weighted quantities does not supply the missing value of the unweighted quotient.

Sources. (4.8)–(4.10), p. 26; (4.11), p. 27. `ProfileHistories.lean:554–588`; `NaturalEntrance.lean:1096–1105`, `NaturalEntrance.lean:1231–1267`; `ActivationStocks.lean:952–970`; `ZerothStressIdentity.lean:334–352`.

4. The pressure bump lies outside the reserved mean interval

Paper: Section 8.3 selects a smooth bump $\hat{\rho}$ with integral one and compact support in $I_m = \{\sqrt{2X} : X \in I_{\text{mean}}\}$. Here I_{mean} is a reserved subinterval of the active annulus; the bump removes the radial integral of the pressure source before compactly supported integration.

Lean: The actual pressure gauge supplies the full annulus endpoints $a = \sqrt{2X_a}$ and $b = \sqrt{2X_b}$ to `PressureStream.rho`. Its centre is $(a + b)/2$ and its outer radius is $(b - a)/4$, so the set on which it is nonzero is

$$\left(\frac{3a + b}{4}, \frac{a + 3b}{4} \right).$$

In the coordinate $\log X$, the reserved mean interval ends more than three units before the outer edge X_b . Passing to $\sqrt{2X}$, its upper endpoint is therefore less than $e^{-3/2}b < b/4$.

Difference: The selected bump is disjoint from the paper's reserved interval, since $(3a + b)/4 > b/4$. Its smoothness and unit integral are proved, and the separate stress-repair bumps use the reserved patch. This establishes a different auxiliary density; it alone does not establish different final pressure or velocity fields.

Sources. §8.3, p. 92. `PressureStream.lean:29–60`; `CorrectionInitialization.lean:4255–4259`; `VariableGaugeMean.lean:58–82`, `VariableGaugeMean.lean:517–523`; `CommonBaseContext.lean:118–125`; `PrimaryTargetBounds.lean:667–668`; `ReservedPatches.lean:29–64`; `NominalConeAssembly.lean:1327–1330`; `OutgoingSchedule.lean:304–306`; `OutgoingTail.lean:542–555`; `CorrectionInitialization.lean:4334–4342`.

5. Spatial phase derivatives pay the time-derivative cost

Paper: In Lemma 9.8, $Q \in (0, 1)$ is a fixed band scale comparable to q , Φ is an oscillatory phase, and $e^{ik\nu\Phi}$ is its carrier, where $k \leq 2Q^{-h/2}$ and the harmonic integer ν is fixed. For a spatial and b time derivatives, with $a + b \geq 1$, the proof retains the respective powers

$$Q^{-a/2-b(1+h)}, \quad Q^{-a(1/2+h/2)-b(1+3h/2)}.$$

These bounds also include constants and fixed powers of $1 + |\log Q|$.

Lean: `localPhasePositiveJets` bounds the full derivative tensor of order $a + b$ using $Q^{-(a+b)(1+h)}$. The carrier and its multiplication by an amplitude are controlled using $Q^{-(a+b)(1+3h/2)}$; the same cost is charged in every direction.

Difference: For $a > 0$, both comparisons lose an additional factor $Q^{-a(1/2+h)}$. These are weaker intermediate bounds, rather than the directional estimates displayed in the paper's proof. Nevertheless, the paper subsequently uses the common loss $\ell_m = 2A + (m+1)(1+3h/2)$, where $A = 1/2 + h$. Lean proves this final loss exactly in `paperPhysicalLoss`; the discrepancy does not concern the conclusion of Lemma 9.8.

Sources. Lemma 9.8 and its proof, pp. 113–114. `SharpCurrentPhysicalPhase.lean:168–175`, `SharpCurrentPhysicalPhase.lean:283–290`; `SharpCurrentModeJets.lean:18–35`; `SharpPhysicalCarrier.lean:12–17`; `WholeDomainPhysicalStageTheorem.lean:238–264`.

6. Frozen scale jets and Euler derivatives give equivalent seminorms

Paper: The normalized profile estimates (5.42) and (5.46) use compositions of $q\partial_q$, ∂_X and ∂_η . Repeating $q\partial_q$ differentiates its coefficient q as well as the profile. The diagonal bound in (5.46) is imposed for $n \geq \max\{2, |I|\}$, where $|I|$ is the derivative order.

Lean: `SlowBorelBase.blownJet` takes ordinary derivatives of $F(q\sigma, X, \eta)$ at $\sigma = 1$, keeping the base scale q fixed. For the pure scale directions, this means

$$\partial_\sigma^2 F(q\sigma, X, \eta)|_{\sigma=1} = q^2 \partial_q^2 F, \quad (q\partial_q)^2 F = q^2 \partial_q^2 F + q\partial_q F.$$

The `AdmissibleScales.blown` condition bounds the tensor of order m of the n th summand by $2^{-n} q^{hn}$ for $n \geq 1$, $m \leq n + 2$, $0 < q \leq 1$, and (X, η) in the chosen compact profile set. The theorem `exists_admissibleScales` constructs a schedule satisfying these conditions.

Difference: The operators are unequal already for $F(q) = q$: the displayed second derivatives are 0 and q . This example compares operators, not the selected profiles. At every fixed order the Euler and frozen derivative families are related by an invertible triangular change of basis with constant coefficients. The same holds for mixed profile derivatives. Consequently they give equivalent all-order $O(q^{2h})$ estimates. Transferring the displayed diagonal bound can introduce an order-dependent multiplicative constant; it does not automatically retain the exact coefficient 2^{-n} . No failure of either the paper's asymptotic estimates or its numerical bound is established here.

Sources. (5.42), p. 60; (5.46), p. 61. `SlowBorelBase.lean:136–139`, `SlowBorelBase.lean:289–355`, `SlowBorelBase.lean:560–566`.

7. The cone margin is chosen on a smaller family

Paper: Appendix C constructs a periodic shear loop with slope $\vartheta(\theta'; \mu)$, mean slope m , and amplitude $\mu \geq 0$. Write $V(\mu) = \langle (\vartheta - m)^2 \rangle$ for its variance, where $\langle \cdot \rangle$ averages over $0 \leq \theta' \leq 2\pi$. The admissible speed satisfies $2 < v < U(\vartheta)$, with U the upper boundary of the stress cone. After choosing $\mu_{\max}$, (C.8) requires

$$U(\vartheta(\theta'; \mu)) > 2 + \delta_L \quad \text{for every } 0 \leq \mu \leq \mu_{\max}, \quad 0 < \delta_L < 1.$$

Lean: `FamilyChoices.cone_margin` instead quantifies over solved tilts with $s = \sqrt{V} \leq \sqrt{3/a_{\min}}$, where $a_{\min} > 0$ is a uniform lower bound for the prescribed angular shear a . Its simultaneously selected `maxAmplitude` satisfies $V(\mu_{\max}) > 3/a_{\min}$, so that endpoint lies outside this family. The selected inequalities are $U \geq 2 + \delta$ and $0 < \delta \leq 1$.

Difference: The selected family is smaller than the full amplitude interval associated with `maxAmplitude`. Its margin suffices for the actual loop, whose strict admissibility is proved. The non-strict auxiliary inequalities are harmless: replacing δ by $\delta/2$ gives a strictly smaller positive margin less than one on the same family. A generic compact-family theorem can also provide a margin on the larger interval, but that is not the selected application. This is a difference in auxiliary choices, not a counterexample to the paper's margin or to the loop construction.

Sources. (C.4)–(C.8), p. 159; (C.9)–(C.10), p. 160. `TrueConeLoop.lean:164–197`, `TrueConeLoop.lean:371–393`, `TrueConeLoop.lean:439–504`; `LoopVariance.lean:842–854`.

8. Uniform full-pressure bounds omit decay of the pressure increment

Paper: Fix the outer and axis data and a sufficiently large reference scale Λ . Let C be the subsequent sufficiently large admissible amplitude parameter, φ_r the normalized reference angular profile, and $\Pi_0(\eta)$ the axis pressure. With $\eta \in [-1, 1]$ and the reference transition width in its prescribed range, the pressure satisfies

$$\Pi_r(X, \eta) - \Pi_0(\eta) = C^{-2} \int_0^X \varphi_r(s, \eta)^2 ds, \quad \sup_{0 \leq X \leq X_i} \|\Pi_r - \Pi_0\|_{C_\eta^k} \leq K_k C^{-2}.$$

Here $X_i = 110$ is fixed and K_k is independent of admissible C and the transition width. Thus (B.24) retains smallness as $C \rightarrow \infty$, at every fixed order.

Lean: `uniform_refF_normalized` controls the η derivatives of Cf , where $f = \varphi_r/C$. The auxiliary theorem `uniform_refF` instead bounds C^{-1} by 1. In `uniform_ref_histories`, this coarser estimate is squared and integrated, and the fixed axis pressure is added. The resulting theorem controls the *full pressure* uniformly in C ; it does not display (B.24)'s C^{-2} bound for the *increment* $\Pi_r - \Pi_0$. Its indicated consumer uses uniform full-pressure control.

Difference: The selected intermediate bound omits the smallness of the pressure increment; a bound for the full pressure is not expected to decay when Π_0 does not. Sharp order-zero and first-parameter-derivative increment bounds are separately proved and used. At every higher fixed order, the normalized-field bounds, the product rule and the pressure integral formula also imply a C^{-2} increment bound. Thus this is a difference in the estimate displayed and used at this step, not evidence that (B.24) fails for the formalized profiles.

Sources. Lemma B.4 and (B.24), pp. 150–151. `ReferenceUniformStocks.lean:245–279`, `ReferenceUniformStocks.lean:285–310`, `ReferenceUniformStocks.lean:341–352`; `ReferenceBounds.lean:96–139`, `ReferenceBounds.lean:955–987`.

9. A full-tensor estimate gives a coarser mixed-derivative bound

Paper: Lemma 10.3 extends the force beyond $t = 1$ by summing $G_j(x, \sigma) = \chi_0(b_j \sigma) \sigma^j F_j(x)/j!$, where $\sigma = t - 1$, F_j is the limiting j th time derivative of the force, χ_0 is a fixed smooth cutoff, and $b_j \geq 1$ controls its width. For a spatial multi-index α and m time derivatives, (10.12) gives

$$\|\partial_x^\alpha \partial_\sigma^m G_j\|_\infty \leq C_{j,m} b_j^{m-j} \|F_j\|_{C^{|\alpha|}}.$$

Only time derivatives differentiate the rescaled cutoff.

Lean: `SpatialBorelExtension.localizedTermDerivBound` estimates the full derivative tensor of total order k on a spatial window by $b_j^{k-j} C_{j,k,W}$. Here $C_{j,k,W}$ is a bound for a fixed template built from F_j on the window W , independent of b_j . Setting $k = |\alpha| + m$ therefore incurs $b_j^{|\alpha|+m-j}$. The proof bounds the norm of the time-rescaling map in all directions.

Difference: Specializing the full-tensor bound to a mixed component gives a coarser scale factor, with an additional $b_j^{|\alpha|}$. This is a cost of the norm estimate, not an actual spatial derivative of the time cutoff. The template constant controls derivatives of F_j up to k , whereas the paper's mixed bound uses only order $|\alpha|$. The pure-time theorem retains b_j^{m-j} , and the coarser joint estimate still proves convergence with every derivative. Spatial windows are used for bounds; their cutoffs do not appear in the final extension series.

Sources. Lemma 10.3, (10.11)–(10.12), p. 120. `SpatialBorelExtension.lean:63–69`,
`SpatialBorelExtension.lean:105–174`, `SpatialBorelExtension.lean:270–304`,
`SpatialBorelExtension.lean:334–388`; `BorelExtension.lean:125–155`.

10. A slower envelope controls the entire future tail

Paper: During the axial transition, the pressure argument uses decay $E(y + s, \eta) \leq K E(y, \eta) e^{-s/2}$ for every $s \geq 0$. Here $y = \log(X/X_R)$ is the logarithmic X -coordinate, with fixed $X_R > 0$, s is an increment in y , and E is the positive unmodified outgoing angular profile. Squaring and integrating this envelope controls the future pressure integral $\Pi(y, \eta) = -\frac{1}{2} \int_y^\infty E(v, \eta)^2 dv$.

Lean: For the actual profile at $\eta = 0$, `logClockFuture` and `finalAngularClockRatio` give

$$E(y + s, 0) \leq e^{3/5} E(y, 0) e^{-2s/5} \quad (y, s \geq 0).$$

The squared decay rate is $4/5$. Comparisons between $E(y, \eta)^2$ and $E(y, 0)^2$ then yield

$$\int_y^\infty E(v, \eta)^2 dv \leq 5e^{6/5} E(y, \eta)^2 \quad (y \geq 0, |\eta| \leq 1).$$

Difference: The selected proof uses the slower rate $2/5$ instead of $1/2$, and transfers a parameter-zero envelope to the physical parameter range. The resulting bound is used for the actual pressure. The original `finalAngularReleaseDecay` proves the sharper rate $1/2 + 3h/4$ after the release stage; its later starting point does not replace the all-future argument used here. No failure of the pressure bound or of tail integrability follows.

Sources. (A.23), p. 133; axial-transition pressure estimate, p. 135; exterior decay, p. 137. `FuturePressureBounds.lean:106–160`,
`FuturePressureBounds.lean:171–217`; `OutgoingEntranceCone.lean:877–889`; `TailCone.lean:227–243`.

11. The local hierarchy is completed before global repair

Paper: Section 5 builds a formal expansion in powers q^{2nh} , $n \geq 0$. Its coefficients are profiles of the velocity and pressure in (X, η) . At order n , Lemma 5.1 solves the equations near the axis using the already finalized lower orders. Lemma 5.2 then extends the current profiles and imposes five integral moment conditions, which confine their stress to an annulus. Only then does the induction proceed to order $n + 1$.

Lean: `SlowRecursion.hierarchy` first recursively constructs every local coefficient, without global moment repair. `schemeFromHierarchy` cuts these local coefficients to produce the seeds for a second recursion. In that global recursion, `exists_step` repairs the current seed and recomputes its pressure source and moments using the already repaired global predecessors.

Difference: The code separates the local and global inductions, whereas the paper interleaves them. The second recursion takes a completed local hierarchy as input. The theorem `profiles_inner_eq` proves agreement of all four velocity/pressure profile components on a common inner region, including the axis, and `profiles_moments` proves the global moment conditions. These are the connecting results needed for the separated construction. This demonstrates a reorganization of the construction and its verification; it does not by itself demonstrate that the resulting coefficients differ from an admissible implementation of the paper's induction.

Sources. §5, p. 45; Lemma 5.1, p. 47; Lemma 5.2, p. 50; Proposition 5.3, p. 55. `SlowRecursion.lean`:931–972; `GlobalSlowProfiles.lean`:770–834, `GlobalSlowProfiles.lean`:1043–1052, `GlobalSlowProfiles.lean`:1090–1118, `GlobalSlowProfiles.lean`:1812–1825.

12. Affine bumps and Rolle's theorem replace geometric dilates

Paper: Lemma 8.7 repairs five moments on an interior radial interval. In the normalized radius $x = r/\sqrt{q}$, the angular moment weights are $x^2, x^{-2-2\lambda}, x^{-2\lambda}$, and the axial weights are $x, x^{1-2\lambda}$, with $\lambda > 0$ fixed. The proof chooses geometric dilates $\beta_j(x) = e^{-jd}\beta_0(e^{-jd}x)$, $j = 0, 1, 2$, of one nonzero nonnegative bump, with $d > 0$. Their moments are $\int x^p \beta_j(x) dx = \mu_p e^{jdp}$, where $\mu_p > 0$. Dividing rows by μ_p gives ordinary Vandermonde matrices. The axial block reuses the first two angular bumps.

Lean: `FiveRowRank` uses affine copies on evenly spaced cells inside (a, b) : for a family of size N , the cell endpoints are $a + (2j + 1)\Delta_N$ and $a + (2j + 2)\Delta_N$, with $\Delta_N = (b - a)/(2N + 1)$ and $0 \leq j < N$. Thus the angular family uses $N = 3$, while the axial family uses the different mesh $N = 2$. Nonsingularity follows by turning a vanishing moment combination into one zero in each cell, then applying Rolle induction to sums of distinct powers after logarithmic change of variable.

Difference: The selected bump families and the proof of invertibility differ from those in Lemma 8.7; in particular, the axial bumps are not the first two angular bumps. However, the paper itself uses Rolle induction for general separated power-moment bumps in Lemma A.1. Thus the formal argument is closely related to another argument already in the paper. The five required moment equations, including their signs, are proved. This is an alternative realization of the moment repair, not missing constraints.

Sources. Lemma 8.7, pp. 97–98; Lemma A.1, pp. 126–127. `FiveRowRank.lean`:24–31, `FiveRowRank.lean`:81–107, `FiveRowRank.lean`:146–156, `FiveRowRank.lean`:240–284; `LocalizedMomentRepair.lean`:21–26, `LocalizedMomentRepair.lean`:87–115; `PowerMomentMatrix.lean`:50–117, `PowerMomentMatrix.lean`:209–249.

DEPARTMENT OF MATHEMATICS, KING'S COLLEGE LONDON

Email address: alexander.bastounis@kcl.ac.uk

DEPARTMENT OF APPLIED MATHEMATICS AND THEORETICAL PHYSICS, UNIVERSITY OF CAMBRIDGE

Email address: fac43@cam.ac.uk

DEPARTMENT OF APPLIED MATHEMATICS AND THEORETICAL PHYSICS, UNIVERSITY OF CAMBRIDGE

Email address: a.hansen@damtp.cam.ac.uk