

Directional failure of moment closures in heterogeneous polar velocity-jump transport

Fujio Takenari

`fujio-takenari-ni@alumni.osaka-u.ac.jp`

Independent Researcher <https://orcid.org/0009-0005-8549-0022>

Research Article

Keywords: velocity-jump process, moment closure, kinetic-to-continuum transport, polar transport, heterogeneous transport, anisotropic diffusion, Monte Carlo simulation, quasi-equilibrium closure, computational biophysics, reproducible computational research, provenance control, fail-closed audit, finite-volume method, oriented cell migration, asymptotic consistency

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11424923/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

Continuum moment closures provide tractable partial differential equations for kinetic descriptions of oriented cell migration, but their finite-scale accuracy can be difficult to separate from numerical or workflow error. We developed a fail-closed, hash-locked audit framework in which independent implementation and adjudication agents evaluate pre-specified derivations, numerical tolerances and execution artifacts under human supervision. We applied this framework to drift-diffusion, anisotropic-diffusion, Chapman-Enskog density and dynamic first-moment quasi-equilibrium closures against a canonical polar velocity-jump Monte Carlo target on a heterogeneous orientation field supported by a real human T1-derived anatomical mask. None of the tested closures recovered adequacy on the frozen final-time metric. For the dynamic closure, the centre-of-mass trajectory diverged directionally from the particle target, with final angular errors of -74.3° and -92.6° in two test cells and combined errors 58–64 times the replicate-noise scale. Removing the finite angular-relaxation initial layer changed the combined error by only 0.34 and 0.23 noise units. These results suggest that local low-order quasi-equilibrium closure can remain inadequate in strongly heterogeneous polar fields despite numerical convergence and asymptotic consistency.

Introduction

Velocity-jump processes provide a foundational kinetic description of biological migration: particles move at fixed speed and reorient stochastically at a Poisson rate toward direction distributions set by their local microenvironment^{1,27–34}. In oriented tissues—including white-matter tracts, collagen networks and aligned extracellular matrix—a common computational strategy is to close the kinetic equation at low angular moments and evolve a small set of continuum fields^{2–9,35}. The resulting drift–diffusion and anisotropic-diffusion equations underpin, among other applications, diffusion-tensor-imaging-informed models of glioma invasion^{7,8,37–41}. The underlying mathematical theory is well developed: diffusion limits are established under parabolic scaling^{2,3,12}, moment hierarchies can be constructed systematically^{10,11}, and first-moment closures that retain flux dynamics have been derived and analysed^{4,13}.

Less attention has been paid to how such closures should be audited when they are used as finite-scale quantitative surrogates for the underlying kinetic process on strongly heterogeneous fields. In that regime, sharp directional changes may arise from the field itself or from processing and modelling conventions, while numerical and workflow choices can introduce errors that are difficult to distinguish from model inadequacy. More generally, computational studies remain vulnerable to post hoc metric selection, in-place repair of failed runs, threshold adjustment and selective reporting of negative outcomes^{18–21,26}. A closure failure is therefore scientifically informative only if alternative explanations arising from the computational pipeline have been constrained in advance.

We address this problem with a fail-closed, dual-agent audit framework for computational model testing. An implementation agent derives, implements, statically audits and numerically qualifies each diagnostic stage inside an isolated, SHA-256-hash-locked namespace. A separate adjudication agent audits the frozen pre-execution package, independently re-derives the governing mathematics, repeats the numerical qualification in a different software environment, and issues or withholds an authorization tied to the approved lock hash. Execution stops on any failed gate; failed criteria and superseded versions are retained rather than overwritten; and acceptance criteria are specified before the tests they judge are first run. A human investigator directs the protocol and brokers all transfers between the two systems.

We then use this framework to test a ladder of local moment closures for a canonical polar velocity-jump process on a heterogeneous orientation field constructed over a real human T1-derived anatomical support mask. The orientation field itself is synthetic, so the study is a controlled transport stress test rather than a biological or clinical validation. The ladder culminates in a dynamic first-moment quasi-equilibrium closure that retains finite flux relaxation and the kinetically correct first-moment initial condition. Its fast-relaxation reduction to the previously audited density closure is verified symbolically with zero residual. Across the frozen final-time metric, none of the tested closures recovers adequacy. Time-resolved analysis of the dynamic closure further reveals a persistent directional mismatch from the Monte Carlo target, while an ablation shows that the finite first-moment relaxation layer contributes only a small fraction of the total discrepancy. Together, these results

provide a traceable test case in which numerical validity and asymptotic consistency are insufficient to guarantee transport-level adequacy.

Results

A provenance-controlled closure ladder

The kinetic reference is a canonical polar velocity-jump process (Methods). Particles move at speed $s = 0.09$ and reorient at rate $\lambda = 10$ toward a local von Mises distribution $q(x, e) \propto \exp[k(x)p(x) \cdot e]$, with concentration $k(x) = 2.5c(x)$ determined by the local orientation coherence $c(x)$ and polar director $p(x) = n(x) \cdot \text{sign}(n \cdot \hat{e})$ under the locked $n \cdot \hat{e} \geq 0$ convention. The structural field is generated once, deterministically, over a real human T1 coronal slice (SHA-256 pinned; dipy sample data²⁵) that defines an anatomical support mask within the unit square; the orientation field itself is synthetic. Two initial-condition cells are considered: a high-coherence seed ('off-singularity', (0.5, 0.45)) and a seed at a nematic defect core ('on-singularity', (0.5, 0.5792...)), the latter re-derived at every run from a locked rule. For each cell, the target statistics are obtained from three Monte Carlo replicates of 30,000 particles with fixed seeds and are required to reproduce seven locked upstream scalar anchors to relative precision 10^{-9} ; the observed maximum relative difference was 0.0.

The frozen primary metric is $E_{\text{frozen}} = E_{\text{RMSE}} + 0.5 \cdot E_{\text{COM}}$, combining the RMSE between normalized final densities with a scaled final centre-of-mass (COM) velocity error ($V_SCALE = 0.04$). Adequacy is reported as E_{frozen} divided by the Monte Carlo replicate-noise scale, which provides the stochastic reference floor for the comparison.

Four diagnostic stages were executed sequentially under the same authorization protocol (Fig. 1): (i) primary model selection among five candidate continuum mechanisms, in which a drift–diffusion model (M3) had the lowest frozen metric among the primary candidates; (ii) an anisotropic-diffusion-plus-polar-drift diagnostic (M34), which did not recover adequacy; (iii) a kinetic-consistent Chapman–Enskog density closure with exact coefficients $u = sA_1(k)p$ and $D = (s^2/\lambda)[(1 - A_2)/2 \cdot I + (A_2 - A_1^2)pp]$, which remained worse than the locked final-time comparators in both cells; and (iv) the dynamic first-moment flux closure examined here. Each stage was hash-locked before the next began, and upstream outputs were used downstream only as read-only evidence.

Dynamic first-moment closure and verified fast-relaxation consistency

The dynamic closure tests whether a density-only representation discards transport information carried by the first angular moment. Taking the first two angular moments of the kinetic equation gives

$$\partial \rho / \partial t + \nabla \cdot J = 0, \quad \partial J / \partial t + \nabla \cdot \Pi = \lambda(\rho u - J),$$

with the second moment closed at quasi-equilibrium, $\Pi_{QE} = \rho(\lambda D + uu) = \rho s^2 E_q[ee]$. Two design properties are central. First, uniform initial headings imply $J(x, 0) = 0$, so the full closure retains the finite first-moment relaxation layer with timescale $\tau = 1/\lambda = 0.1$. An ablation initialized with $J(x, 0) = \rho_0 u$ ('instant relaxation') isolates the effect of this first-moment initial layer. This does not make the complete initial angular state exact: where $c > 0$, the quasi-equilibrium second moment Π_{QE} differs from the true uniform-heading second moment at $t = 0$. Second, adiabatic elimination of J at large λ recovers the audited density closure, $J = \rho u - \nabla \cdot (Dp) - (\rho/\lambda)(u \cdot \nabla)u + O(\lambda^{-2})$. This reduction is re-derived symbolically inside the static audit at every run, and the observed residual is zero in both components.

The solver treats (ρ, J, J_y) as a hyperbolic relaxation system with a Rusanov finite-volume flux^{15,42,45} and an integrator that is exact for the stiff relaxation term in the homogeneous benchmark (error $\leq 10^{-17}$)^{16,17,43}. Ten pre-specified qualification tests all pass. These include asymptotic recovery of the density closure under a criterion fixed before its first run (monotone residual shrink of 64.5× over $\lambda = 10$ to 300) and a production-horizon time-convergence gate showing that the frozen time resolution contributes 2.55% of the smaller replicate-noise scale to the metric-form discretization error. The numerical time-step budget was therefore fixed below the stochastic reference floor before the scientific run.

The qualification history also exposed two classes of implementation risk. An early pipeline version contained a genuine integrator defect that was invisible on homogeneous tests but was detected by a heterogeneous regression identity; the corrected production path is now bound to that identity at execution time. Separately, two pre-specified asymptotic-recovery criteria failed at coarse time resolution. Those failures were retained, independently reproduced in a second software environment, and diagnosed by controlled step-size refinement as time-discretization error. The time step was then refined under unchanged thresholds, while the superseded pipeline versions remained de-authorized. Qualification values subsequently agreed across environments to full printed precision.

Structural audits define the locked transport challenge

Two mandatory pre-execution audits characterize the locked field before the closure comparison (Fig. 2). Applying the $n \geq 0$ polarity convention to the smooth synthetic director field generates 520 sign-discontinuity edges, with a maximum drift jump of 0.1376, approximately twice the maximum local drift magnitude, across one grid cell. Both seeds lie one grid cell (0.0099 length units) from the nearest seam, but in very different local environments: the off-singularity seed is adjacent to a high-coherence seam ($c = 0.961$; jump 0.1357), whereas the on-singularity seed lies near the low-coherence defect core ($c = 0.025$; jump 0.0026). A separate boundary audit shows that the actual outer boundary of the locked A1 field is isotropic with zero drift, so the adopted reflecting ghost-cell treatment is exact for this specific boundary configuration. These audits diagnose the geometry; they do not smooth or repair it.

The dynamic closure does not recover adequacy in either cell

The authorized production run lasted 109.8 s. All provenance stages and hard numerical-validity gates passed: relative mass change was $\leq 3.9 \times 10^{-16}$, no clipping occurred, and the boundary mass-flux residual was 0 (Table 2). The run produced four records, {off, on} $\times$ {full closure, instant-relaxation ablation}, each with the full time-resolved state stored at the ten mandatory times and independently re-verified from the persisted arrays by a read-only validator.

Table 1. Main results for the four dynamic-closure records on the frozen Monte Carlo-anchored metric.

Record	E_density	E_COM	E_comb.	Adequacy ($\times$ noise)	vs primary	vs diagnostic	vs density closure	Final angle err.	Final speed (closure/MC)
off / full closure	0.11632	0.68953	0.46109	63.72	-28.9%	-29.3%	-0.68%	-74.32°	0.01031 / 0.02852
off / instant ablation	0.11597	0.68533	0.45864	63.38	-28.2%	-28.6%	-0.15%	-73.31°	0.01047 / 0.02852
on / full closure	0.09908	0.57279	0.38548	57.80	-12.7%	-16.2%	-1.97%	-92.62°	0.00731 / 0.02138
on / instant ablation	0.09858	0.57074	0.38395	57.58	-12.2%	-15.7%	-1.56%	-91.75°	0.00738 / 0.02138

Negative percentages indicate a larger (worse) dynamic-closure error than the named comparator. Replicate noise: 0.00724 (off-cell) and 0.00667 (on-cell). Initial-layer ablation effect (full closure minus instant-relaxation ablation): +0.00245 (0.34 noise units, off-cell) and +0.00153 (0.23 noise units, on-cell).

Table 2. Pre-specified numerical-validity gates; all four dynamic-closure records satisfied every gate.

Gate	Tolerance	Worst observed value
Finiteness throughout	required	all finite
Clip ratio	$\leq 10^{-6}$	0.0
Mass conservation (relative)	$\leq 10^{-6}$	3.9×10^{-16}
Boundary mass-flux residual	$\leq 10^{-10}$	0.0
Normalized flux bound	≤ 1.05	0.752
Maximum time step vs. bound	step $\leq$ bound	$4.1265 \times 10^{-4} \leq 4.1267 \times 10^{-4}$
Substep-vs-production regression	= 0.0	0.0

The final-time outcome is consistent across both cells (Table 1; Figs. 3–5). For the off-singularity cell, the full dynamic closure yields $E_{\text{dynamic}} = 0.4611$ (63.7× replicate noise), which is 28.9% larger than the locked drift–diffusion primary, 29.3% larger than the anisotropic-diffusion diagnostic and 0.7% larger than the density closure. For the on-singularity cell, the corresponding value is 0.3855 (57.8× replicate noise), 12.7%, 16.2% and 2.0% larger than the same three comparators. Removing the first-moment relaxation layer changes the combined error by only + 0.00245 (off) and + 0.00153 (on), equivalent to 0.34 and 0.23 replicate-noise units. Thus, within this test, the finite first-moment initial layer has a small effect on the final discrepancy.

Time-resolved analysis identifies a persistent directional mismatch

Time-resolved comparison of the dynamic closure with the regenerated Monte Carlo target localizes the dominant discrepancy (Fig. 4). Through $t = 0.2$ (2τ), the COM separation remains ≤ 0.006 length units. Thereafter, the trajectories diverge: the Monte Carlo COM moves predominantly in + y, with final headings of + 93.5° (off) and + 92.0° (on), whereas the full dynamic closure ends at + 19.2° and – 0.6°, respectively. The resulting final angular errors are – 74.3° and – 92.6°, and the dynamic-closure COM speeds are lower than the Monte Carlo values by factors of 2.8 and 2.9. The full and instant-relaxation variants remain close to one another after the early transient, while their error relative to the Monte Carlo target continues to grow. This separation indicates that the first-moment initial layer is not the dominant source of the late-time transport mismatch.

The integrity record constrains alternative computational explanations

Across four audit rounds, the framework detected a silent implementation defect that homogeneous tests had missed, retained two pre-specified criterion failures without threshold adjustment, rejected superseded lock hashes even when paired with valid authorization tokens, and reproduced qualification values across software environments to full printed precision. The reported observables are contained in, or recomputable from, the hash-verified handoff package, and the persisted state arrays reproduce the corresponding derived quantities to 10^{-12} or exactly. This audit record does not establish biological validity, but it substantially narrows the class of numerical and provenance failures that could account for the observed closure discrepancy.

Discussion

For the canonical polar velocity-jump process and locked heterogeneous field studied here, none of the tested local continuum closures recovers adequacy on the frozen final-time metric. The time-resolved evidence is strongest for the dynamic first-moment closure: despite numerical convergence, exact conservation checks and verified fast-relaxation consistency with the density closure, its centre-of-mass trajectory diverges markedly in direction from the Monte Carlo reference in both test cells. The instant-relaxation ablation further shows that the finite first-moment initial layer contributes only a small fraction of the final discrepancy. These observations support a narrower conclusion than a general rejection of

moment methods: under this strongly heterogeneous polar field, local low-order quasi-equilibrium closure is insufficient to reproduce the reference transport trajectory.

This finding does not conflict with rigorous diffusion-limit results for velocity-jump processes^{2,3,12}; rather, it highlights the distinction between an asymptotic limit and finite-scale quantitative adequacy. Here, the locked polarity convention creates a dense network of drift discontinuities, including a high-coherence near-reversal seam adjacent to one seed. Such discontinuities violate the smooth-field setting underlying many asymptotic arguments and can amplify differences between local continuum sampling and particle transport. The present evidence does not identify the unique missing mechanism. It instead motivates controlled tests of higher angular moments, nonlocal transport descriptions, or explicit treatment of seam structure¹⁰. The dynamic first-moment result also cautions against assuming that adding hyperbolic relaxation alone will restore adequacy when a diffusion-type closure fails^{4,13}.

The audit framework provides a complementary methodological contribution. Pre-registration and reproducibility norms^{18–21,54–57} emphasize the value of prespecified criteria and traceable computational workflows; here those principles are implemented through separated implementation and adjudication roles, hash-pinned artifacts, fail-closed execution and preserved failure histories. The value of this architecture lies not in AI autonomy but in enforced separation of roles and human-accountable verification. As AI-assisted research workflows become more common^{22,46–49}, comparable provenance controls may help prevent rapid iteration from obscuring criterion changes, superseded results or implementation drift.

Several limitations constrain the interpretation. The anatomical content is a single real human T1 coronal slice used only as a support mask, with a synthetic orientation field superimposed; this is a controlled geometry stress test rather than a biological or clinical validation. Only two seed locations, one field and one kinetic parameter set were examined. The seam network is induced by the locked polarity convention, so a different convention would define a different transport problem. The Monte Carlo target is a regenerated, identity-gated surrogate for the upstream target rather than a reused raw array, and time-resolved comparisons are therefore available only for the dynamic closure versus the regenerated Monte Carlo target; upstream closures enter only as final-time read-only comparators. Finally, although $J(x,0) = 0$ matches the uniform-heading first moment exactly, the quasi-equilibrium second moment does not match the true uniform-heading second moment at $t = 0$ in regions with $c > 0$. The present study therefore motivates, but does not perform, the next step beyond a first-moment closure.

Methods

Canonical kinetic model and structural field

The kinetic reference is a canonical polar velocity-jump model taken from five source files whose SHA-256 hashes are locked and re-verified before every run. Particles at position $x \in [0,1]^2$ with heading $e \in S^1$ satisfy $\partial f / \partial t + s e \cdot \nabla f = \lambda \{ \rho(x,t) q(x,e) - f \}$, with $\rho = \int_{S^1} f de$, $s = 0.09$ and $\lambda = 10$. The post-turn kernel is $q(x,e) = \exp[k(x)p(x) \cdot e] / (2\pi I_0(k(x)))$, with $k = 2.5c(x)$. The polar director is $p = n \cdot \text{sign}(n \cdot \nabla)$, with $\text{sign}(0) \rightarrow +1$. The structural field $(\theta(x), c(x))$ is generated deterministically: a real human T1 coronal slice from the dipy sample data²⁵ (SHA-256 verified at runtime; 256×256) is thresholded and resampled to an anatomical support mask on an $N = 101$ cell-centred grid, over which a synthetic coronal orientation field is constructed. The support mask is used only for post hoc leakage diagnostics and is never treated as a solver boundary. The two seeds are (0.5, 0.45) and (0.5, 0.5792079207920792); the latter is re-derived at every run from the locked minimum-coherence rule.

Monte Carlo target and identity gate

For each cell, the target is regenerated once as three replicates of 30,000 particles using fixed seeds, an MC step of 0.005 and horizon $T = 6$. Final densities are Gaussian-smoothed and normalized exactly as in the upstream workflow. Seven scalar anchors—replicate noise for both metric components, two turning-kernel order parameters, turning-rate ratio, COM speed and mask leakage—must match the locked upstream evidence to relative tolerance 10^{-9} ; the observed maximum relative

difference was 0.0 in both cells. The Monte Carlo COM trajectory was additionally sampled at the ten mandatory times 0, 0.025, 0.05, 0.1, 0.2, 0.5, 1, 2, 4 and 6.

The dynamic closure system, discretization, and frozen time resolution

The first- and second-moment definitions $J = \int e f de$ and $\Pi = \int e^2 f e de$ yield the moment system in Results by direct integration over e . The quasi-equilibrium closure uses $s^2 E_q[ee] = \lambda D + uu$, with $u = sA_1(k)p$, $A_j(k) = I_j(k)/I_0(k)^{14,36}$, and D as defined in Results. No fitted parameter is introduced in this closure. The full variant uses $J(\cdot, 0) = 0$, matching the first moment of the uniform initial heading distribution; the instant-relaxation ablation uses $J(\cdot, 0) = \rho_0 u$. The full variant should not be interpreted as matching the complete initial angular state: Π_{QE} differs from the true uniform-heading second moment at $t = 0$ where $c > 0$. Reflecting boundaries are implemented with ghost cells (density Neumann; antisymmetric normal flux; symmetric tangential flux). For the actual locked A1 outer boundary, $c = 0$ and the drift vanishes, so this treatment coincides with the corresponding specular-reflection moment condition.

The solver advances (ρ, J_x, J_y) with a Rusanov/local-Lax–Friedrichs finite-volume flux^{15,42,44,45}, using local wave speeds $\sqrt{M}$ and $\sqrt{M_{yy}}$ for $M = \lambda D + uu$, together with an integrator that is exact for the homogeneous relaxation term^{16,17,43}. A single authoritative substep is used by the production loop and is linked to it by a bit-exact regression gate (tolerance 0.0) in both qualification and execution. The production resolution is fixed at CFL = 0.003125; the maximum step used was 4.1265×10^{-4} against a bound of 4.1267×10^{-4} , corresponding to approximately 1.5×10^4 substeps per record. This resolution was selected solely from a pre-run synthetic time-convergence gate, for which the metric-form discretization difference at the final refinement was 1.698×10^{-4} , or 2.55% of the smaller locked replicate-noise scale, with a step-halving ratio of 2.01.

Qualification suite, criterion history, and hard gates

Ten mandatory pre-execution tests, run only on synthetic or manufactured fields, assess coefficient identities, homogeneous relaxation (error $\leq 10^{-17}$), finite-relaxation pulse kinematics, asymptotic recovery of the density closure, mass conservation, boundary residual, grid and time convergence, heterogeneous stability, a manufactured polarity-seam stress test, and the substep/production regression identity. The asymptotic-recovery criterion history is retained in full: two pre-specified criteria failed at coarse time resolution, the failure was independently reproduced and diagnosed by controlled refinement, and the final configuration passed under unchanged thresholds with monotone residual shrink of $64.5\times$. Hard execution gates require finiteness, clip ratio $\leq 10^{-6}$, relative mass change $\leq 10^{-6}$, boundary residual $\leq 10^{-10}$, a recorded production CFL, and an exact substep/production regression identity. A global normalized-flux diagnostic, $\max |J|/(s \cdot \max p_{\text{initial}}) \leq 1.05$, is used as a numerical sanity check (worst observed 0.752) and is not interpreted as a pointwise realizability proof. Any failed gate halts execution without automatic repair or rerun.

Dual-agent provenance protocol and state artifacts

Scientific execution requires both an explicit authorization flag and an approved SHA-256 lock hash matching a nine-file pre-execution lock over the execution-critical source and configuration files. Three superseded candidate locks were verified to be rejected even when paired with a valid authorization flag. A sixteen-stage fail-closed harness checks the pre-execution lock, canonical source hashes, cleanliness of the working directory, software-environment identity, the input-data hash, the upstream-evidence lock, and a live static audit before execution. Post-execution stages verify successful completion, the presence and structure of the result ledger and all four state artifacts, an independent read-only validator, and a final SHA-256 lock generated after re-reading and freshly re-hashing the written files. For each record, the raw state at the ten mandatory times (density and both flux components, full floating-point precision) is stored without downsampling or precision reduction. The validator independently reproduces the stored density hashes, final densities, COM trace and flux diagnostics and re-checks the hard numerical gates. Time-resolved comparison is restricted to the dynamic closure versus the regenerated Monte Carlo target; upstream comparators are used only at final time because their raw time series were not serialized.

Roles, adjudication, and statistics

Anthropic Claude served as the implementation agent for derivation support, implementation, qualification and provenance-controlled execution under human authorization. OpenAI ChatGPT served as the independent adjudication agent for mathematical re-derivation, cross-environment qualification checks, execution authorization and post-execution scientific audit. The human investigator specified the protocol, brokered transfers between the systems, determined the study design and retained responsibility for all scientific decisions and reporting. Uncertainty is summarized by the Monte Carlo replicate-noise scale, defined as the mean pairwise distance among the three replicates on the combined metric (0.00724 off-cell; 0.00667 on-cell). Adequacy ratios and noise-unit effect sizes are reported rather than null-hypothesis significance tests, and closure-adequacy rulings use sign-based comparisons against locked final-time comparators. No fitting, parameter tuning, or post hoc change to the frozen metric or adjudication threshold was introduced during the reported stage.

Declarations

Competing interests

The authors declare no competing interests.

Ethics declaration

Ethics declaration: not applicable. This is a purely computational study involving no human or animal subjects, no collection of primary data from participants, and no clinical material. The anatomical geometry is derived from a standard, publicly available sample dataset (dipy T1 sample data), and all orientation fields are synthetic; no ethics approval or participant consent was therefore required.

AI usage statement

Two AI systems were used as documented research tools under continuous human direction. Anthropic Claude supported mathematical derivation, implementation, numerical qualification, provenance-controlled execution and factual reporting; OpenAI ChatGPT supported independent mathematical checks, cross-environment qualification review, execution authorization and post-execution scientific audit. All AI-assisted outputs entering the scientific record were subject to human review and are traceable within the deposited provenance package. Neither AI system is an author; the corresponding human author retains responsibility for the accuracy and integrity of the manuscript.

Funding

This research received no external funding.

Author contributions

Takenari Fujio conceived the study, specified and enforced the dual-agent audit protocol, brokered transfers between the two AI systems, supervised the diagnostic stages, and drafted and revised the manuscript. **Hiromi Fujio**: Independent Researcher, Ono, Hyogo, Japan. Provided critical administrative and logistical support for protocol coordination, data archiving, and independent cross-validation oversight throughout the multi-agent audit process. Both authors approved the submitted version and accept responsibility for their respective contributions. AI-system roles are described in Methods and in the AI usage statement; no AI system is listed as an author.

Data availability

The provenance package underlying this study—including the result ledger, four raw state arrays at all ten mandatory times, cryptographic locks and manifests, static-audit and numerical-qualification reports, adjudication records, and canonical source files—will be deposited in a public repository. The T1 input is the standard dipy sample dataset²⁵.

Code availability

All execution-critical code used for closure evaluation, numerical integration, orchestration, qualification, static audit, independent validation and provenance control is included in the same repository and is sufficient to reproduce the reported workflow in the stated software environment. Figures were generated with Matplotlib⁵⁸.

References

1. Othmer HG, Dunbar SR, Alt W (1988) Models of dispersal in biological systems. *J Math Biol* 26:263–298
2. Hillen T, Othmer HG (2000) The diffusion limit of transport equations derived from velocity-jump processes. *SIAM J Appl Math* 61:751–775
3. Othmer HG, Hillen T (2002) The diffusion limit of transport equations II: chemotaxis equations. *SIAM J Appl Math* 62:1222–1250
4. Hillen T (2004) On the L^2 -moment closure of transport equations: the Cattaneo approximation. *Discrete Contin Dyn Syst B* 4:961–982
5. Hillen T (2006) M^5 mesoscopic and macroscopic models for mesenchymal motion. *J Math Biol* 53:585–616
6. Chauviere A, Hillen T, Preziosi L (2007) Modeling cell movement in anisotropic and heterogeneous network tissues. *Netw Heterog Media* 2:333–357
7. Painter KJ, Hillen T (2013) Mathematical modelling of glioma growth: the use of diffusion tensor imaging (DTI) data to predict the anisotropic pathways of cancer invasion. *J Theor Biol* 323:25–39
8. Engwer C, Hillen T, Knappitsch M, Surulescu C (2015) Glioma follow white matter tracts: a multiscale DTI-based model. *J Math Biol* 71:551–582
9. Hillen T, Painter KJ (2011) Transport and anisotropic diffusion models for movement in oriented habitats. In *Dispersal, Individual Movement and Spatial Ecology* Lecture Notes in Mathematics 177–222 (Springer, 2013)
10. Levermore CD (1996) Moment closure hierarchies for kinetic theories. *J Stat Phys* 83:1021–1065
11. Grad H (1949) On the kinetic theory of rarefied gases. *Commun Pure Appl Math* 2:331–407
12. Chalub FACC, Markowich PA, Perthame B, Schmeiser C (2004) Kinetic models for chemotaxis and their drift-diffusion limits. *Monatsh Math* 142:123–141
13. Joseph DD, Preziosi L (1989) Heat waves. *Rev Mod Phys* 61:41–73
14. Mardia KV, Jupp P (2000) *E. Directional Statistics* (Wiley)
15. Toro EF (2009) *Riemann Solvers and Numerical Methods for Fluid Dynamics*, 3rd edn. Springer
16. Pareschi L, Russo G (2005) Implicit–explicit Runge–Kutta schemes and applications to hyperbolic systems with relaxation. *J Sci Comput* 25:129–155
17. Jin S (1999) Efficient asymptotic-preserving (AP) schemes for some multiscale kinetic equations. *SIAM J Sci Comput* 21:441–454
18. Peng RD (2011) Reproducible research in computational science. *Science* 334:1226–1227
19. Sandve GK, Nekrutenko A, Taylor J, Hovig E (2013) Ten simple rules for reproducible computational research. *PLoS Comput Biol* 9:e1003285

20. Nosek BA, Ebersole CR, DeHaven AC, Mellor DT (2018) The preregistration revolution. *Proc. Natl Acad. Sci. USA* 115, 2600–2606
21. National Academies of (2019) Sciences, Engineering, and Medicine. Reproducibility and Replicability in Science. National Academies
22. Wang H et al (2023) Scientific discovery in the age of artificial intelligence. *Nature* 620:47–60
23. Harris CR et al (2020) Array programming with NumPy. *Nature* 585:357–362
24. Virtanen P et al (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* 17:261–272
25. Garyfallidis E et al (2014) Dipy, a library for the analysis of diffusion MRI data. *Front Neuroinform* 8:8
26. Ioannidis JPA (2005) Why most published research findings are false. *PLoS Med* 2:e124
27. Patlak CS (1953) Random walk with persistence and external bias. *Bull Math Biophys* 15:311–338
28. Keller EF, Segel LA (1970) Initiation of slime mold aggregation viewed as an instability. *J Theor Biol* 26:399–415
29. Alt W (1980) Biased random walk models for chemotaxis and related diffusion approximations. *J Math Biol* 9:147–177
30. Berg HC, Brown DA (1972) Chemotaxis in *Escherichia coli* analysed by three-dimensional tracking. *Nature* 239:500–504
31. Erban R, Othmer HG (2004) From individual to collective behavior in bacterial chemotaxis. *SIAM J Appl Math* 65:361–391
32. Dolak Y, Schmeiser C (2005) Kinetic models for chemotaxis: hydrodynamic limits and spatio-temporal mechanisms. *J Math Biol* 51:595–615
33. Codling EA, Hill NA (2005) Calculating spatial statistics for velocity jump processes with experimentally observed reorientation parameters. *J Math Biol* 51:527–556
34. Codling EA, Plank MJ, Benhamou S (2008) Random walk models in biology. *J R Soc Interface* 5:813–834
35. Hillen T, Painter K (2009) J. A user's guide to PDE models for chemotaxis. *J Math Biol* 58:183–217
36. Fisher NI (1993) *Statistical Analysis of Circular Data*. Cambridge Univ. Press
37. Basser PJ, Mattiello J, LeBihan D (1994) MR diffusion tensor spectroscopy and imaging. *Biophys J* 66:259–267
38. Jbabdi S et al (2005) Simulation of anisotropic growth of low-grade gliomas using diffusion tensor imaging. *Magn Reson Med* 54:616–624
39. Swanson KR, Alvord EC Jr, Murray JD (2000) A quantitative model for differential motility of gliomas in grey and white matter. *Cell Prolif* 33:317–329
40. Swanson KR, Bridge C, Murray JD, Alvord EC (2003) Jr. Virtual and real brain tumors: using mathematical modeling to quantify glioma growth and invasion. *J Neurol Sci* 216:1–10
41. Alfonso JCL et al (2017) The biology and mathematical modelling of glioma invasion: a review. *J R Soc Interface* 14:20170490
42. Jin S, Xin Z (1995) The relaxation schemes for systems of conservation laws in arbitrary space dimensions. *Commun Pure Appl Math* 48:235–276
43. Filbet F, Jin S (2010) A class of asymptotic preserving schemes for kinetic equations and related problems with stiff sources. *J Comput Phys* 229:7625–7648
44. LeVeque RJ (2002) *Finite Volume Methods for Hyperbolic Problems*. Cambridge Univ. Press
45. Roe PL (1981) Approximate Riemann solvers, parameter vectors, and difference schemes. *J Comput Phys* 43:357–372
46. Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* 378:686–707
47. Jumper J et al (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596:583–589
48. Davies A et al (2021) Advancing mathematics by guiding human intuition with AI. *Nature* 600:70–74
49. Trinh TH, Wu Y, Le QV, He H, Luong T (2024) Solving olympiad geometry without human demonstrations. *Nature* 625:476–482
50. Marchetti MC et al (2013) Hydrodynamics of soft active matter. *Rev Mod Phys* 85:1143–1189

51. Tailleur J, Cates ME (2008) Statistical mechanics of interacting run-and-tumble bacteria. *Phys Rev Lett* 100:218103
52. Cates ME (2012) Diffusive transport without detailed balance in motile bacteria: does microbiology need statistical physics? *Rep. Prog Phys* 75:042601
53. Alert R, Trepas X (2020) Physical models of collective cell migration. *Annu Rev Condens Matter Phys* 11:77–101
54. Baker M (2016) 1,500 scientists lift the lid on reproducibility. *Nature* 533:452–454
55. Munafò MR et al (2017) A manifesto for reproducible science. *Nat Hum Behav* 1:0021
56. Open Science Collaboration (2015) Estimating the reproducibility of psychological science. *Science* 349:aac4716
57. Stodden V et al (2016) Enhancing reproducibility for computational methods. *Science* 354:1240–1241
58. Hunter JD (2007) Matplotlib: a 2D graphics environment. *Comput Sci Eng* 9:90–95

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [fig1fieldandliftReference.png](#)
- [fig2seamcontrolReference.png](#)
- [fig3diagnosisReference.png](#)
- [fig4consistentReference.png](#)