Return to phase space

with an AI assistant

Anjor Kanekar · Vienna · 5 August 2026

Ten years out

2015 — PhD, Maryland

2019 — Meyrand et al., PNAS echo paper

2026 — back

A decade in industry

AI lowered the barrier to re-entry making a return to physics possible in 2026.

Dissipation Anomaly as a test problem

Fluids — known: as viscosity → 0, turbulent dissipation does not vanish; the cascade delivers energy to small scales at a fixed rate.

Phase space — expected, but (afaik) never verified numerically: the kinetic cascade should do the same in velocity space — ε_ν independent of ν.

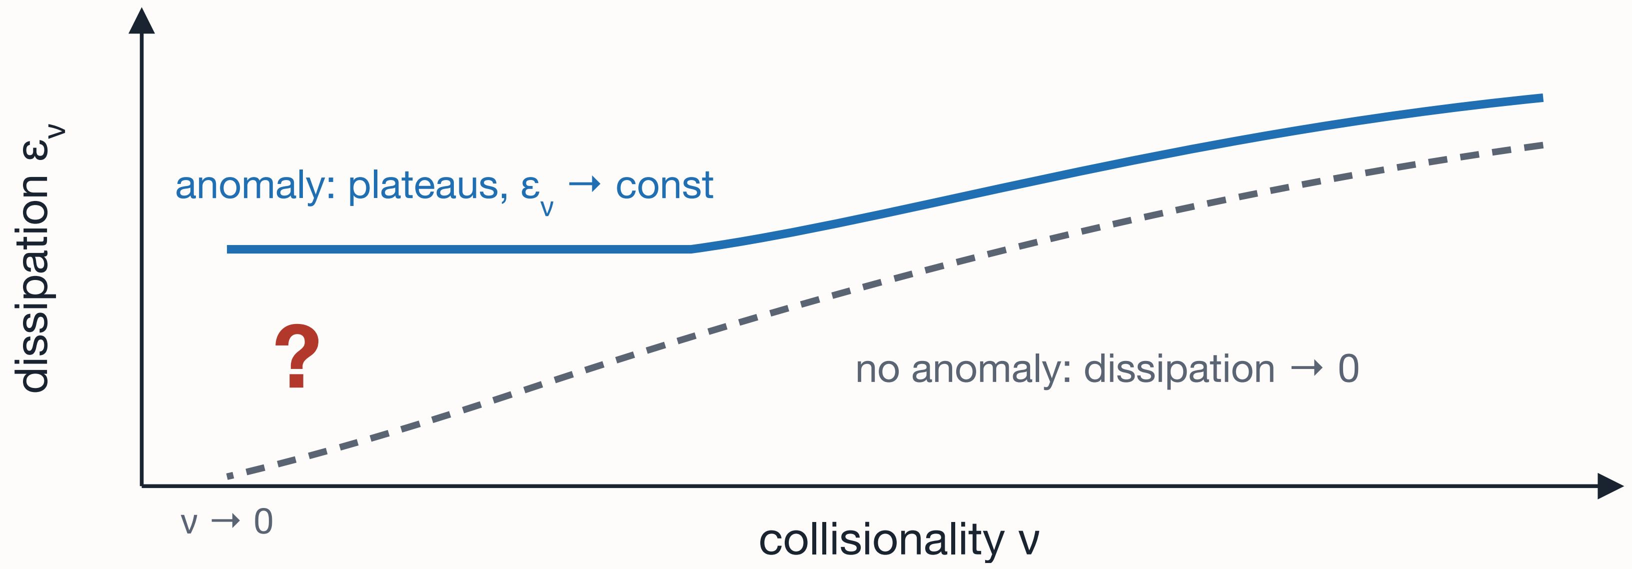

This was a good first target to test with AI - it has an answer that's easy to grade, but still a certain level of complexity.

What I (well, Claude) built

GAndAlf 2.0 — JAX KRMHD spectral solver · Fourier × Hermite · pip-installable · same code, MacBook M1 → Modal A100

Solves inertial-range KRMHD — Alfvén waves governed by RMHD; compressive fluctuations (δn, $\delta B_\parallel$) passively mixed by the Alfvénic cascade, via a kinetic equation for g — Schekochihin et al. 2009, §§ 5–6

“**100%** of this code was written by Claude Code.

I have read almost none of it.”

github.com/anjor/gandalf

How you trust code you haven’t read

1 · **Alfvén-wave dispersion** — linear physics, exact answer known	day one
2 · **Orszag–Tang vortex** — nonlinear fluid benchmark	week one
3 · **Spectral checks** — $k_\perp^{-5/3}$ turbulent cascade; Hermite $m^{-1/2}$ phase mixing vs. theory	escalating
4 · **Energy conservation + balance gates** — automated, every run, **failures never reach analysis**	every production run

You don’t trust it by reading it. You make it reproduce known physics at escalating difficulty, and validate the physics.

Code review is replaced by physics review.

The run log

docs/run_log.md — 32 logged runs. 26 never reached analysis.

run	hardware	outcome	gates
01_M004_20260310	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ✗
02_nu1e-3_20260325	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ✗
02_nu1e-2_20260325	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ✗
02_nu1e-3_20260326	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ok
02_nu1e-2_20260326	MacBook	FAIL	conservation ✗ · balance ✗ · steady ok · spectral ✗
⁝ 19 more runs, 25–31 March – all FAIL ⁝			
02_nu1e-3_lowdrive_20260330	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ok
02_alfven_bench_20260331	MacBook	FAIL	conservation ✗ · balance ✗ · steady ✗ · spectral ok
02_hermite128_nu3_imex_20260417	Modal A100	pass	$\varepsilon_\nu = 49.21 \pm 10.65$ · all gates green
02_hermite128_nu5_imex_20260418	Modal A100	pass	$\varepsilon_\nu = 49.26 \pm 10.61$ · matches $\nu=3$ to 0.1%
02_hermite128_nu10_imex_20260418	Modal A100	pass	$\varepsilon_\nu = 49.23 \pm 10.81$ · matches $\nu=3$ to 0.1%
02_hermite128_nu1_imex_20260418	Modal A100	pass	$\varepsilon_\nu = 48.42 \pm 11.41$ · plateau extends to $\nu=1$
02_hermite128_nu20_imex_20260418	Modal A100	pass	$\varepsilon_\nu = 49.21 \pm 10.99$ · plateau holds at high ν
02_hermite128_nu50_imex_20260418	Modal A100	pass	$\varepsilon_\nu = 49.21 \pm 11.21$ · plateau across 50× in ν

Every red line is a run that never reached analysis.

Results

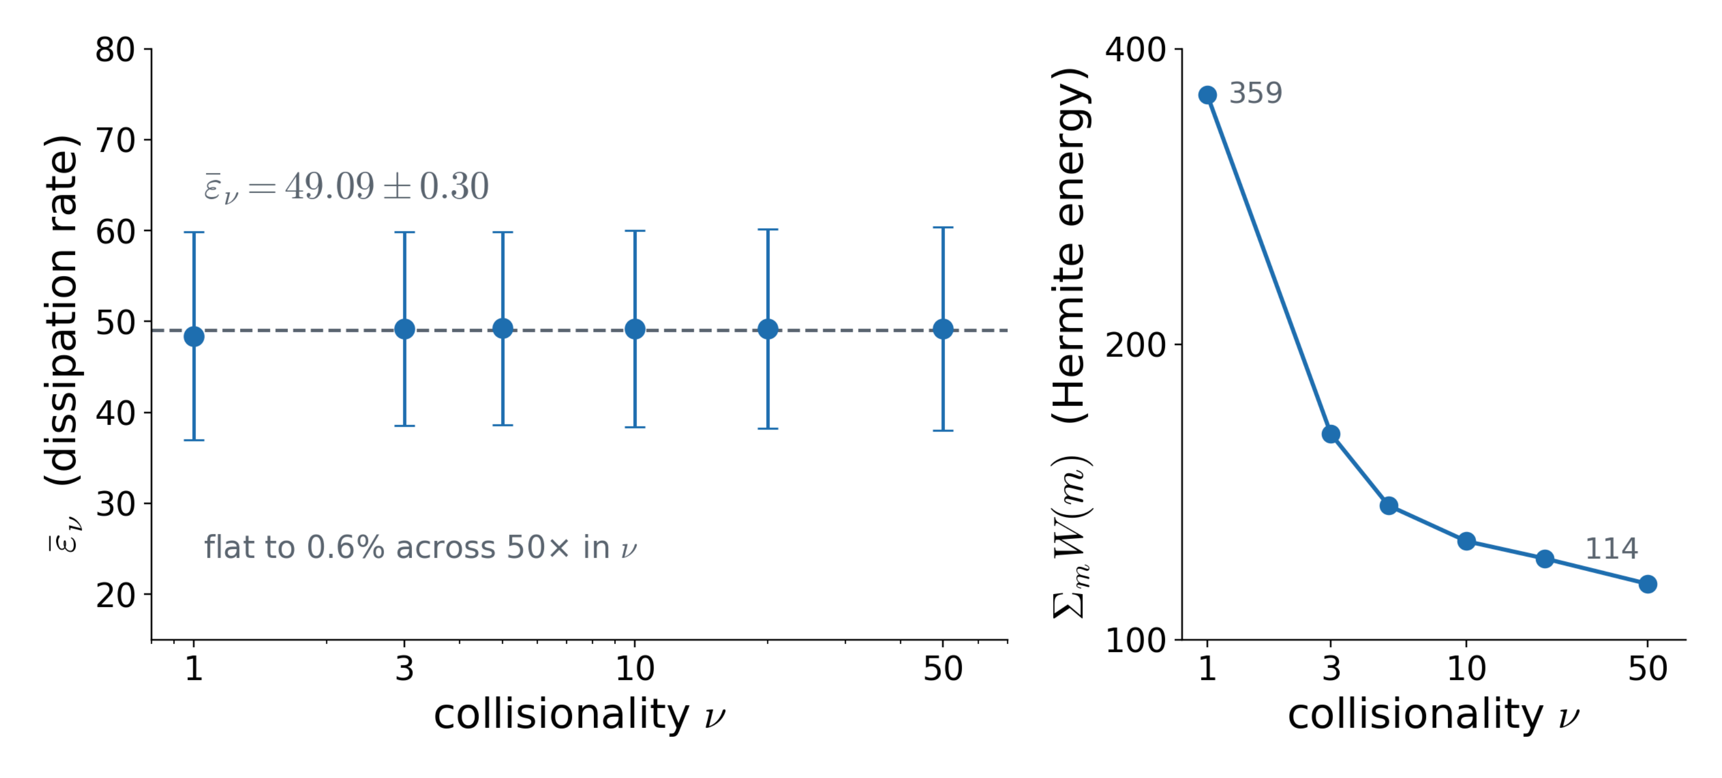

M = 128 Hermite moments · 128^3 · 200 τ_A per run · all gates green. Dissipation rate constant to 0.6% across 50× in collisionality, while the velocity-space cascade extends to higher m as ν drops (total Hermite energy 114 → 359 as ν goes 50 → 1) — **the expected anomaly**.

A deliberately unsurprising result.

The economics

1
person, ten years rusty

9
weeks

~$150
of GPU time

41.5 A100-hours logged for the production v-scan (6 runs × ≈7 h, Modal) · everything else ran on a MacBook

Off-the-shelf tools. No scaffolding team. No HPC allocation.

Archis and Thomas showed you the high end on Monday — I’m the lower bound. If I can do this, you certainly can.

What the machine could not do

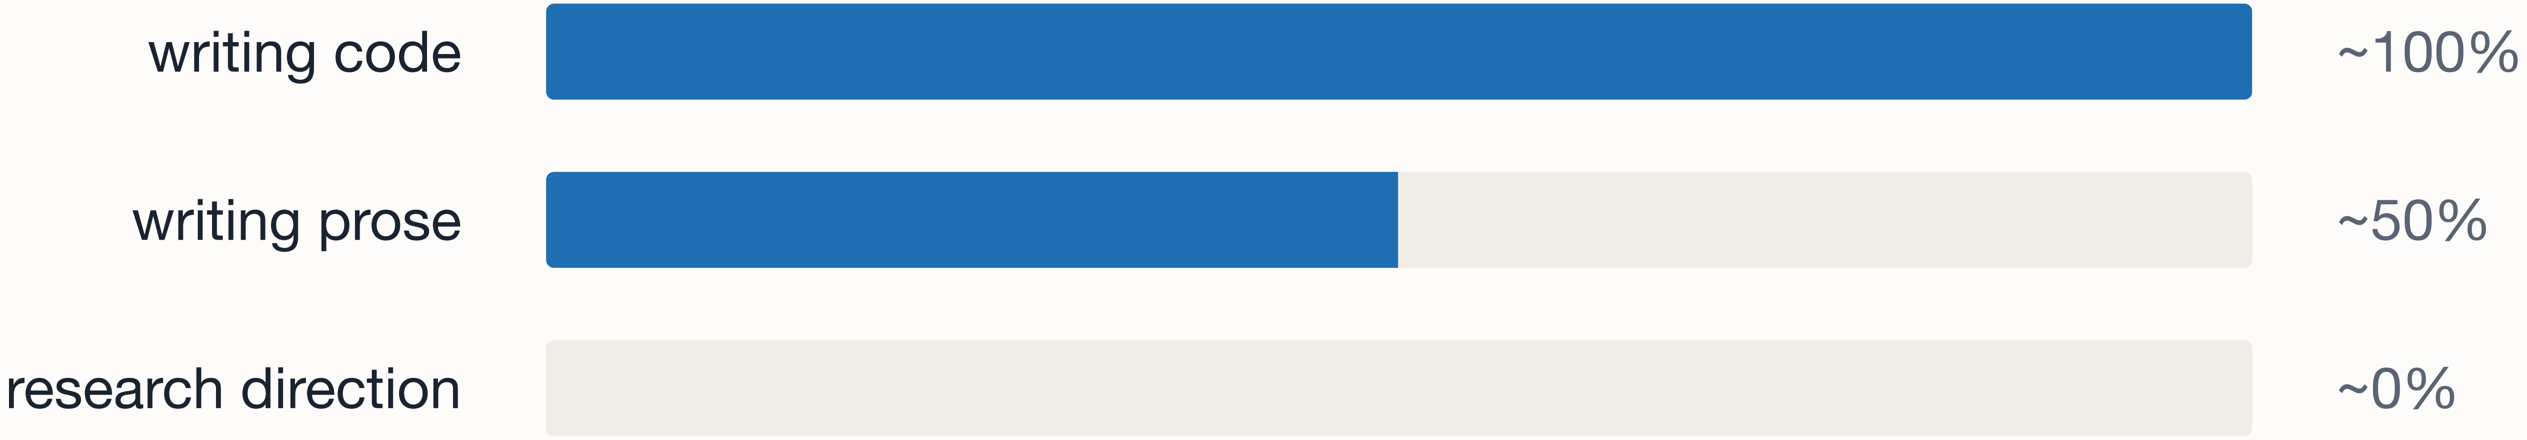

- No physics taste — jumps toward answers instead of exploring parameter space methodically
- Cannot tell when a result *looks wrong*
- Every consequential choice — the question, the gates, when to stop trusting a run — was human

It does the part you’re bored by. The part you’re good at is still yours.

What's the version of this for *your* problem?

Thursday, AI discussion (with Thomas & Archis) — hands-on.
Bring a concrete problem!

Backup slides

Backup · Hermite spectra W(m) vs ν

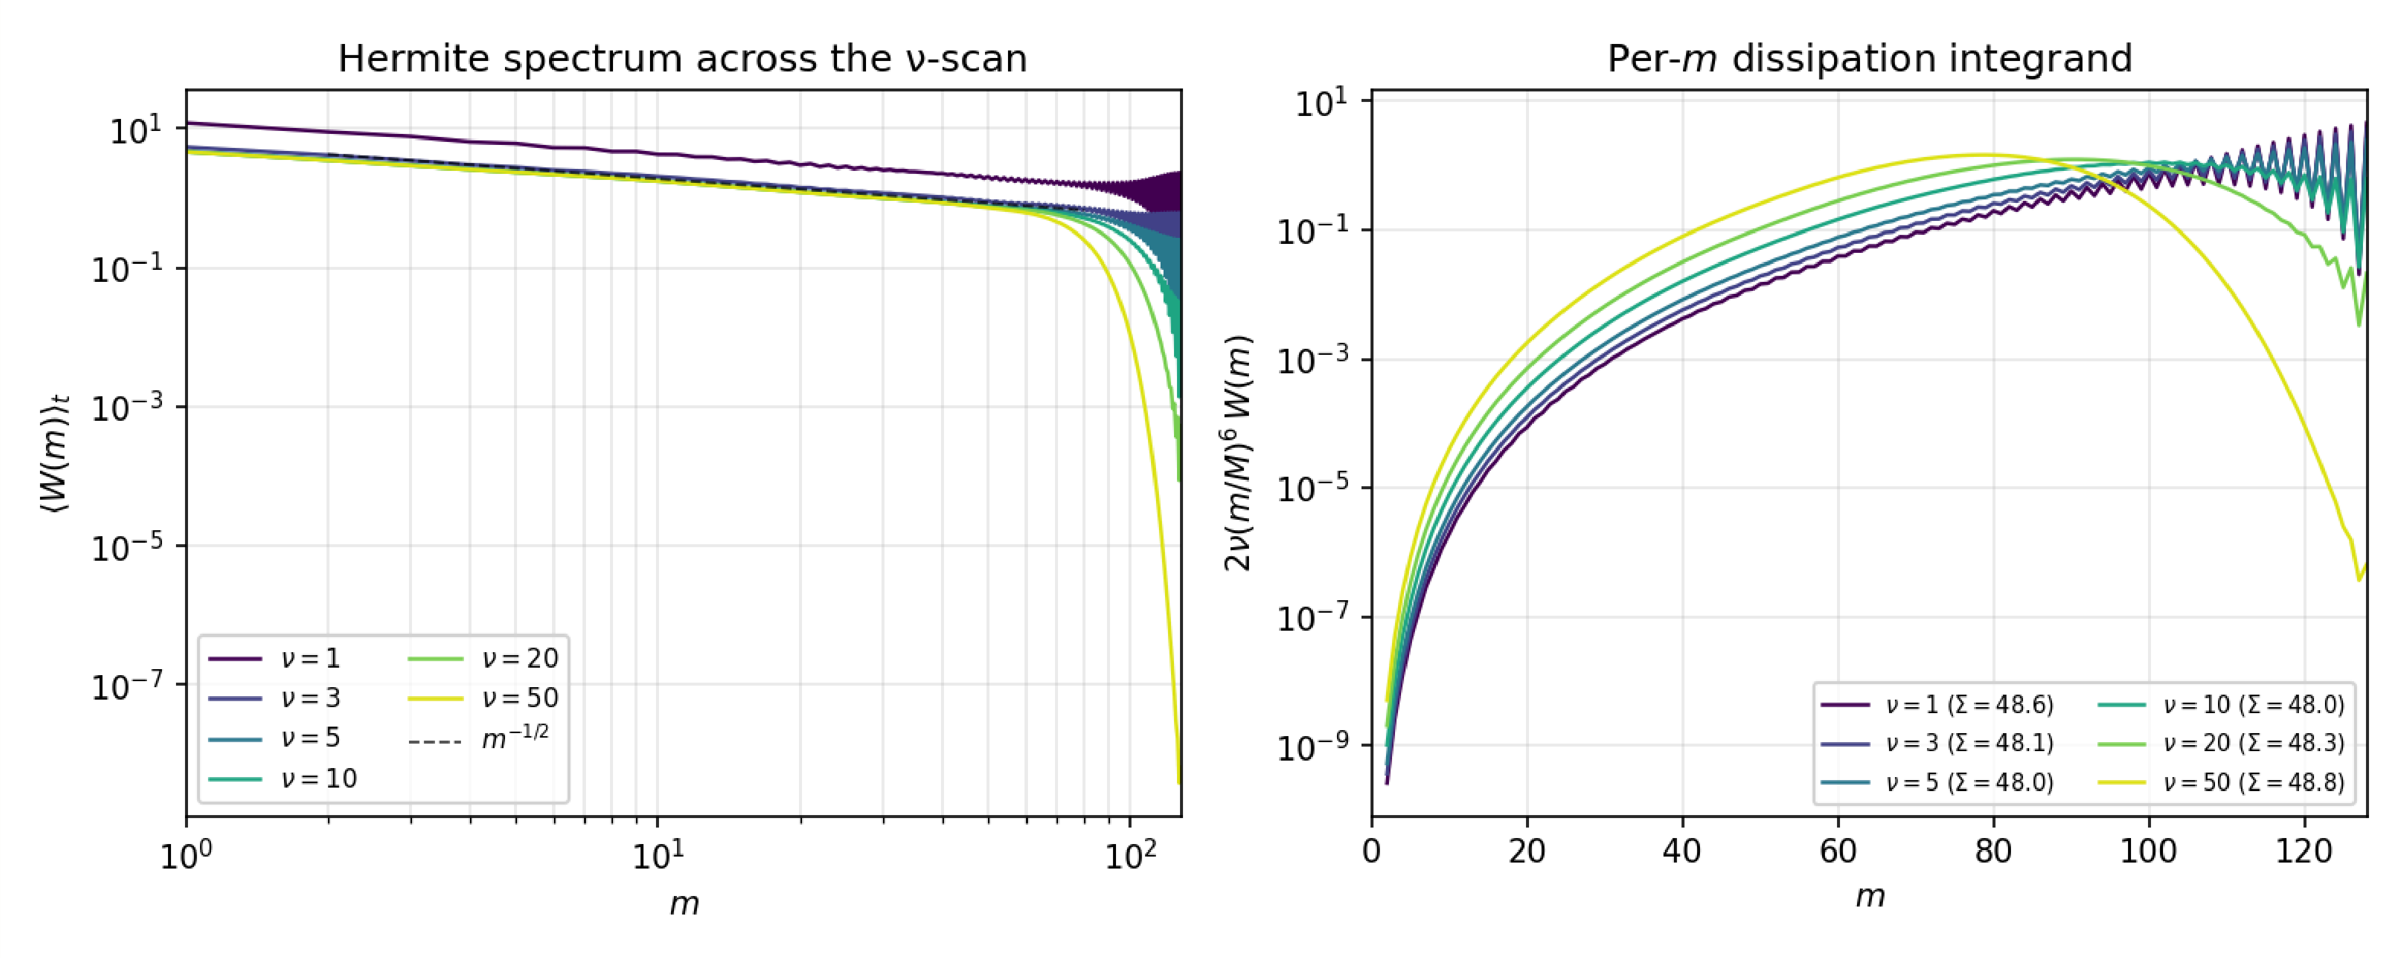

Inertial-range slope within 5% of the Zocco–Schekochihin $m^{-1/2}$ phase-mixing prediction across all six ν (–0.472 … –0.496). ΣW(m) self-adjusts 359 → 114 with a kink near ν ≈ 3.

Backup · ν = 3 production time series

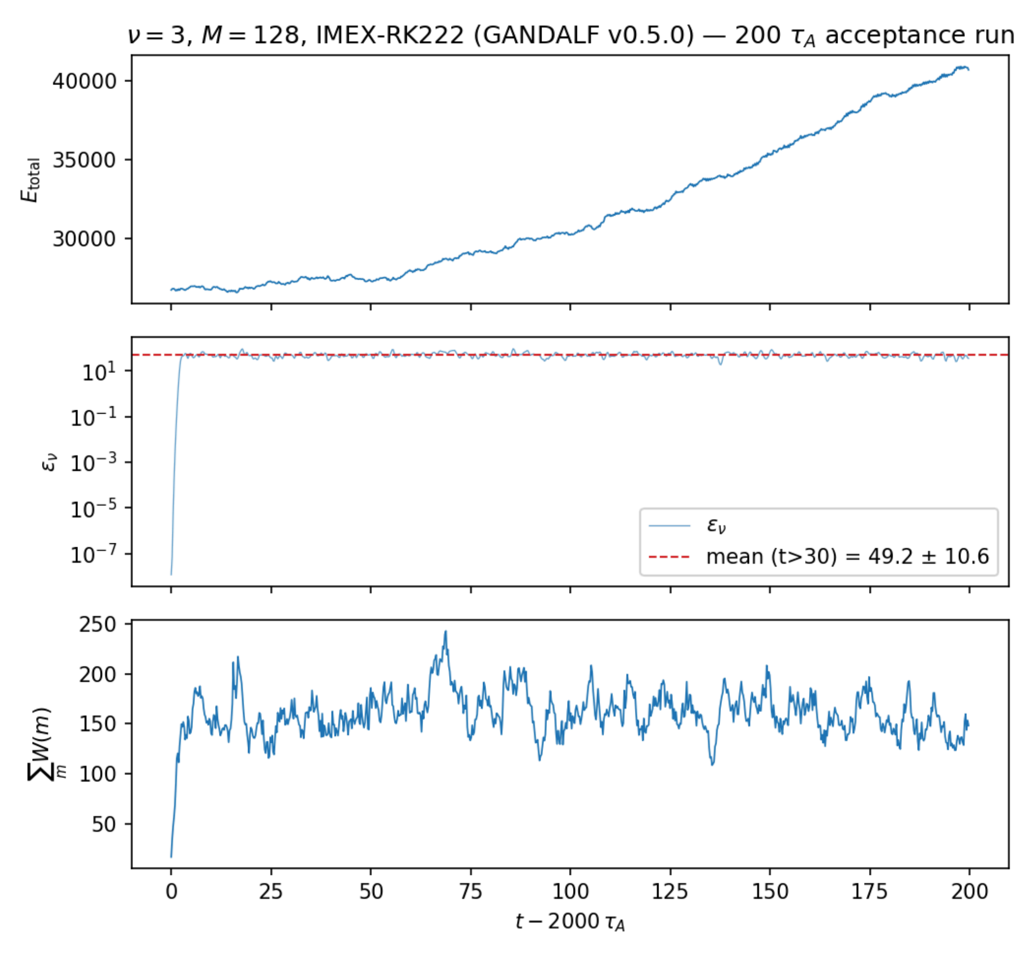

200 τ_A from the Alfvénic checkpoint at t = 2000. ε_ν steady after a ~30 τ_A Hermite fill-in transient; E_total drifts ×1.52 (filling in) — identical across all ν.

Backup · Convergence & resolution

- **Why M = 128:** M is the phase-space resolution whose cascade is under study — lowering it masks the physics rather than testing it
- Companion runs at M = 64 and M = 256 (ν = 3) to confirm resolution
- Hyper-collisionality $(m/M)^n$: m = 0, 1 exempt; dissipation confined to high m at low ν
- Spatial: 128^3, forcing at low k; $E(k_\perp)$ shows a clean inertial range (spectral-sanity gate)

Backup · Claude.md

```
# CLAUDE.md — Critical Rules (verbatim)

## 2. Physics validation before science
Every simulation must pass these gates before extracting any scientific result:
- Energy conservation: relative error < 1% at steady state
- Energy balance: injection rate ≈ dissipation rate (within 5%)
- Steady state: total energy fluctuations < 10% of mean over final 50 τ_A
- Spectral sanity: E(k⊥) shows inertial range

If a gate fails, **stop and report**. Do not proceed to analysis.

## 4. Run identification
Every run gets a unique ID ... Log every run to docs/run_log.md with:
ID, config path, hardware, wall time, outcome (pass/fail), notes.
```

The validation gates are literally instructions to the agent — the run log you saw is their output.

Backup · The March graveyard — what was actually wrong

- **Symptom:** every nonlinear Hermite run failed gates or blew up (26 red rows, 10–31 March)
- **Diagnosis:** explicit (Lawson-RK4) treatment of the hyper-collision term is numerically unstable at $M = 128$ — the 50 τ_A “results” were transients of an unstable scheme, not physics
- **Fix:** IMEX-RK222 — implicit collision term (GANDALF v0.5.0, issue #137)
- **Then:** all six production runs pass every gate, first try

The gates caught a subtle numerical instability that plausible-looking output would have smuggled past a human reader.

Backup · One concrete loop

1. **Prompt:** "ε_ν drifts upward without saturating in the $\nu = 10^{-3}$ run — diagnose against the energy-balance gate."
2. **Agent:** reads diagnostics, plots k_z spectra, isolates the $z\pm$ high-k_z growth, files the analysis
3. **Gate check:** proposed fix re-run against all four gates on a smoke test
4. **Human:** decides the result still looks wrong → escalate to scheme change, not parameter tweak

The human contribution is step 4 — and choosing which prompts are worth asking.