WHAT LIES BEYOND THE CUT

How much turbulence can a coarse grid know about?

Approaching the optimal closure: equivariance, inductive bias, and Reynolds-number generalization in data-driven LES

Syver Døving Agdestein · Benjamin Sanderse

ECCOMAS 2026 · MS184A Theory-guided Design of Deep Learning-based Surrogates

→ advance · ← back · print → PDF backup

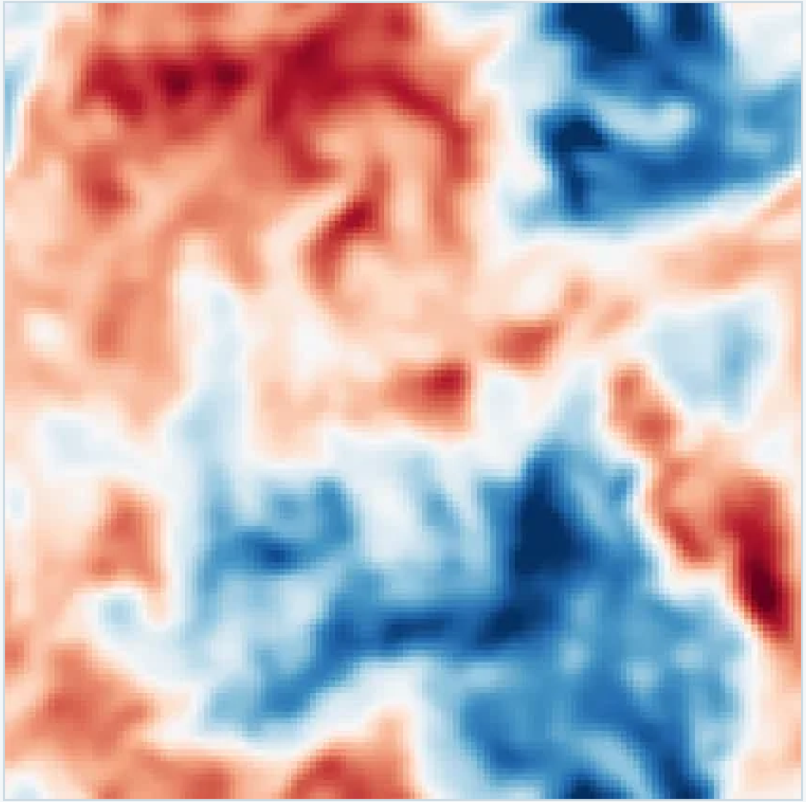

what an LES grid sees of a turbulent flow

The same flow, on two grids

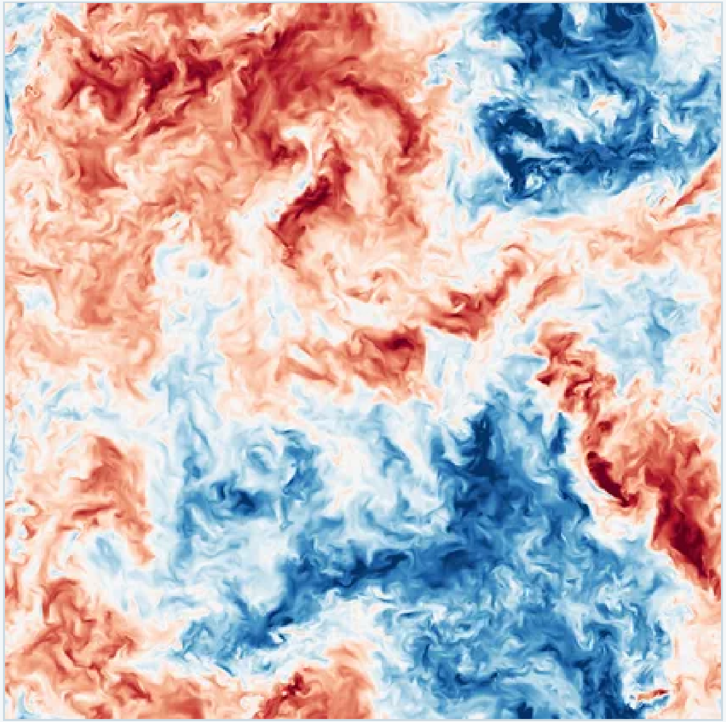

DNS · VELOCITY u

810^3 — resolves everything

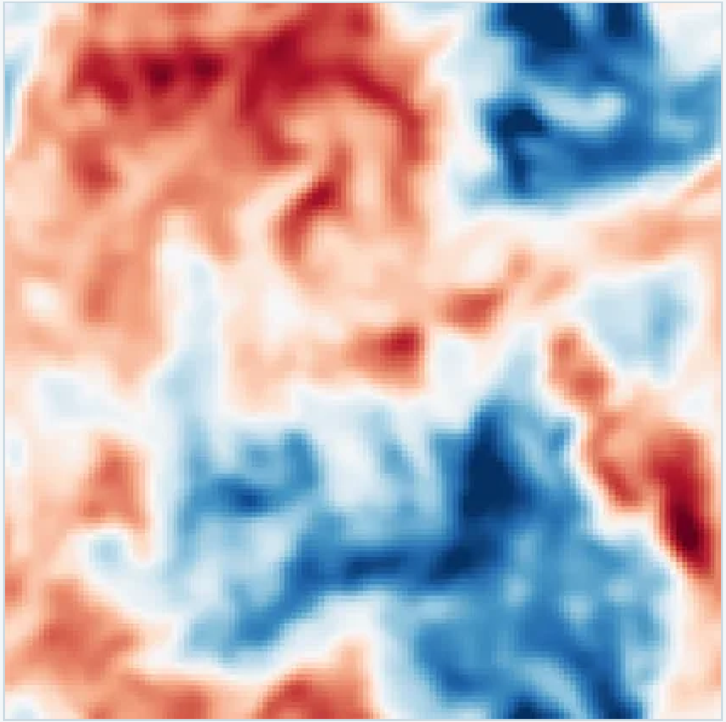

LES · FILTERED VELOCITY ū

128^3 — 250× fewer points, misses the small scales

the missing scales still act on the resolved ones — that action is the closure's job

Where the energy lives — and where the grid stops

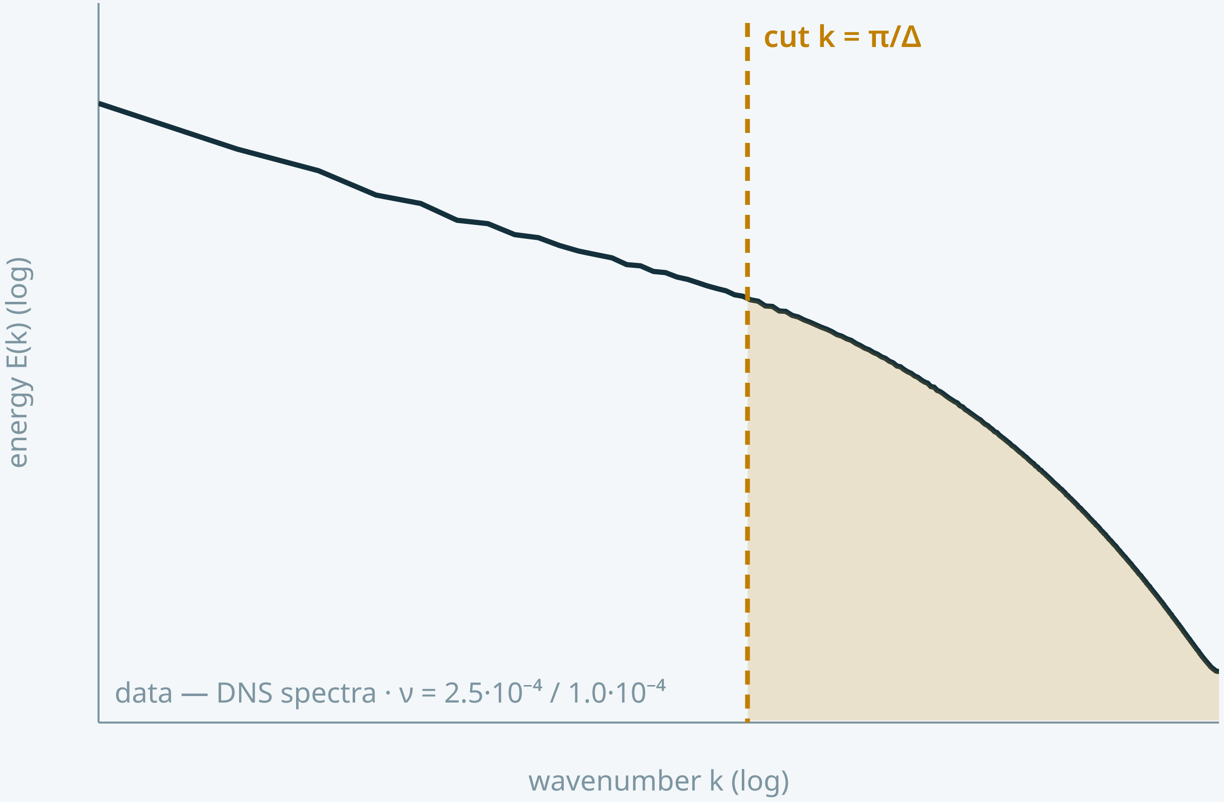

the grid resolves everything left of the cut;
the shaded energy is invisible

coarsen the filter (Δ ↑) — more energy hidden

raise the Reynolds number (ν ↓) — the spectrum grows a longer tail

$$\mathrm{Re}_\Delta = \frac{\Delta^2 \, \|\bar{A}\|}{\nu}$$

one resolved-scale number for
how much lies beyond the cut

Why this number? Scale the equations and see what survives

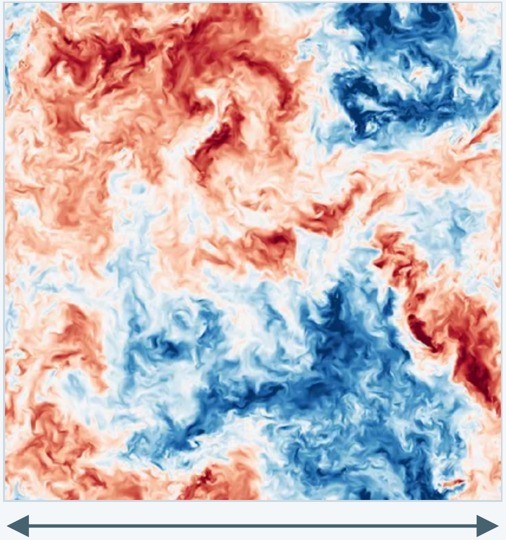

a turbulent solution

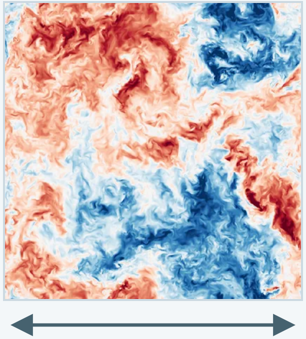

…rescaled – still a solution

$$x \to ax, \quad t \to bt, \quad u \to \frac{a}{b}\,u, \quad \nu \to \frac{a^2}{b}\,\nu$$

Navier–Stokes maps to Navier–Stokes — a two-parameter symmetry.
The fixed LES grid keeps only the **$a^2 = b$** family.

along that family, exactly one dimensionless group built from ū is invariant:
Re_Δ — the symmetry analysis hands the closure its Reynolds input

Energy crosses the cut in both directions

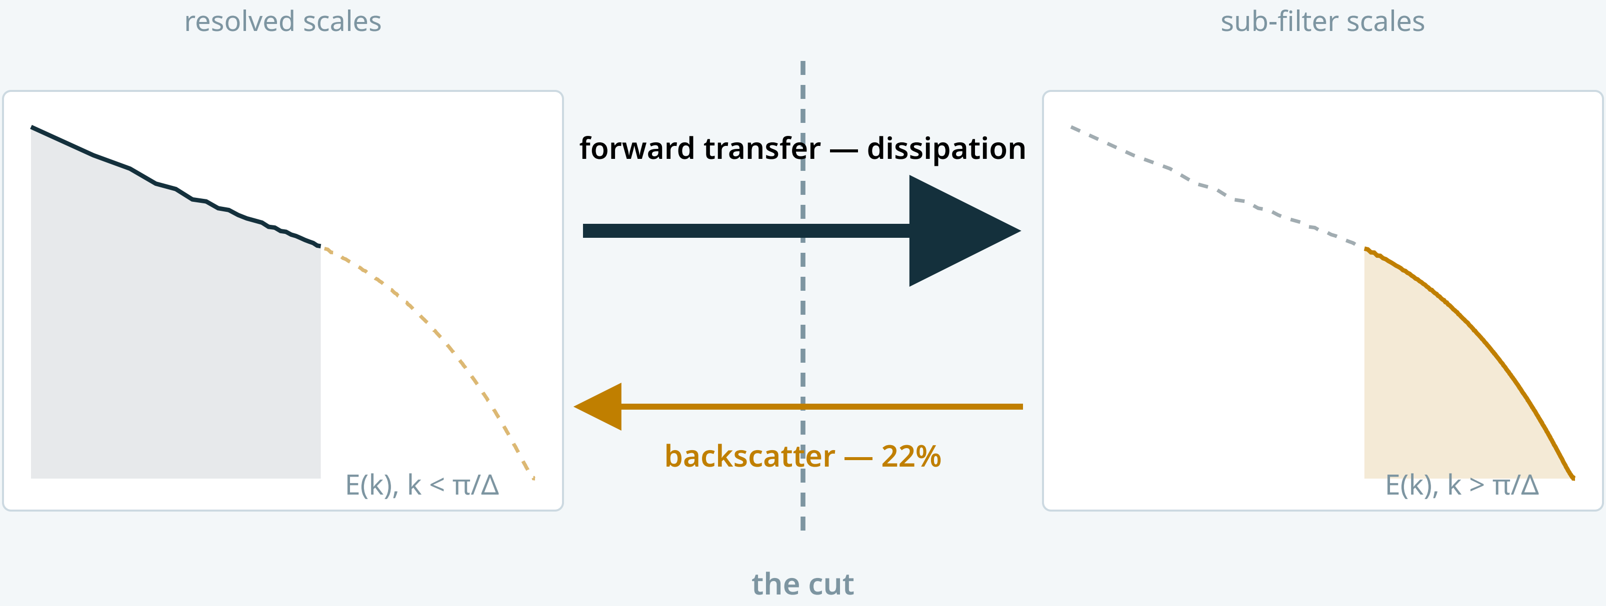

$\Pi = -\tau : \bar{S}$ a closure must get this net flux right — at every Reynolds number

Classical closures pick one side of this trade

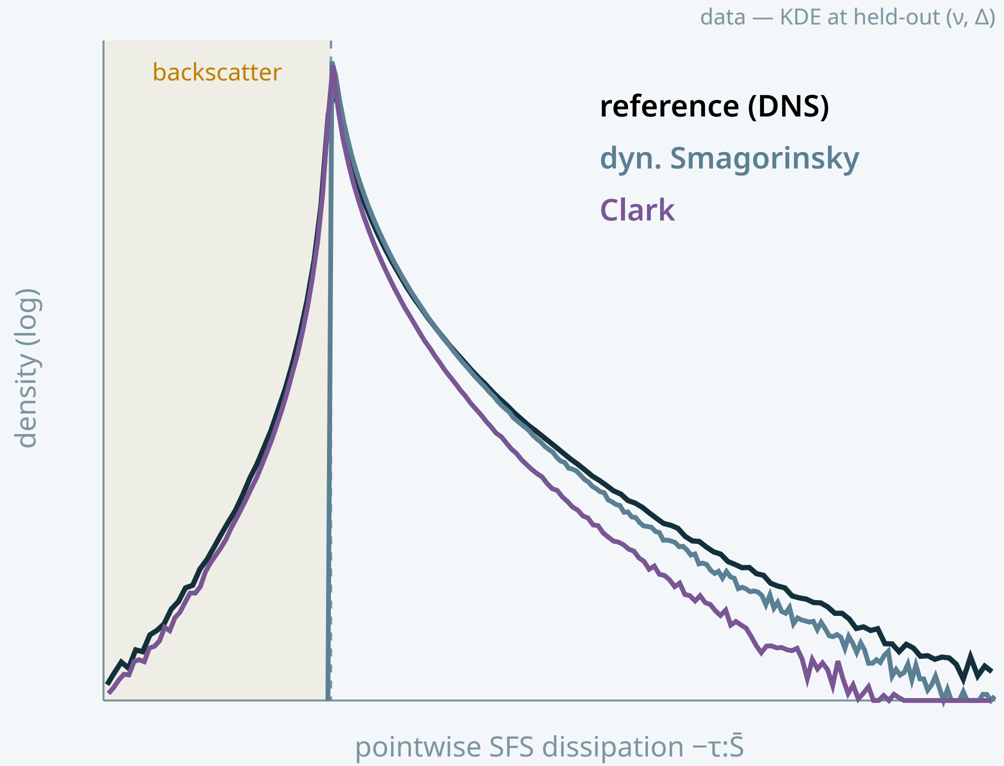

FUNCTIONAL · SMAGORINSKY

$$\tau = -2\,(C_s\Delta)^2\,|\bar{S}|\,\bar{S}$$

dissipation ≈ right · structure wrong

`backscatter 0.0004 — can't flow backward`

STRUCTURAL · CLARK

$$\tau = \frac{\Delta^2}{12}\,\bar{A}\,\bar{A}^{\mathsf{T}}$$

structure ≈ right · dissipation 0.6× too weak

`backscatter 0.25 — close to the true 0.22`

and notice: **ν appears in neither formula** no ν

The learned closures are just as blind

$$\tau = \Delta^2 |\bar{A}|^2 \, m\left(\bar{A}/|\bar{A}|\right)$$ no ν

the normalization that grants scale-invariance also **erases the Reynolds number** — the dissipation calibration freezes at the training regime

fix: train across viscosities, and give every model Re_Δ as one extra input

$$\tau = \Delta^2 |\bar{A}|^2 \, m\left(\bar{A}/|\bar{A}|,\ \mathrm{Re}_\Delta\right)$$ with ν

Tested away from training: who keeps their calibration?

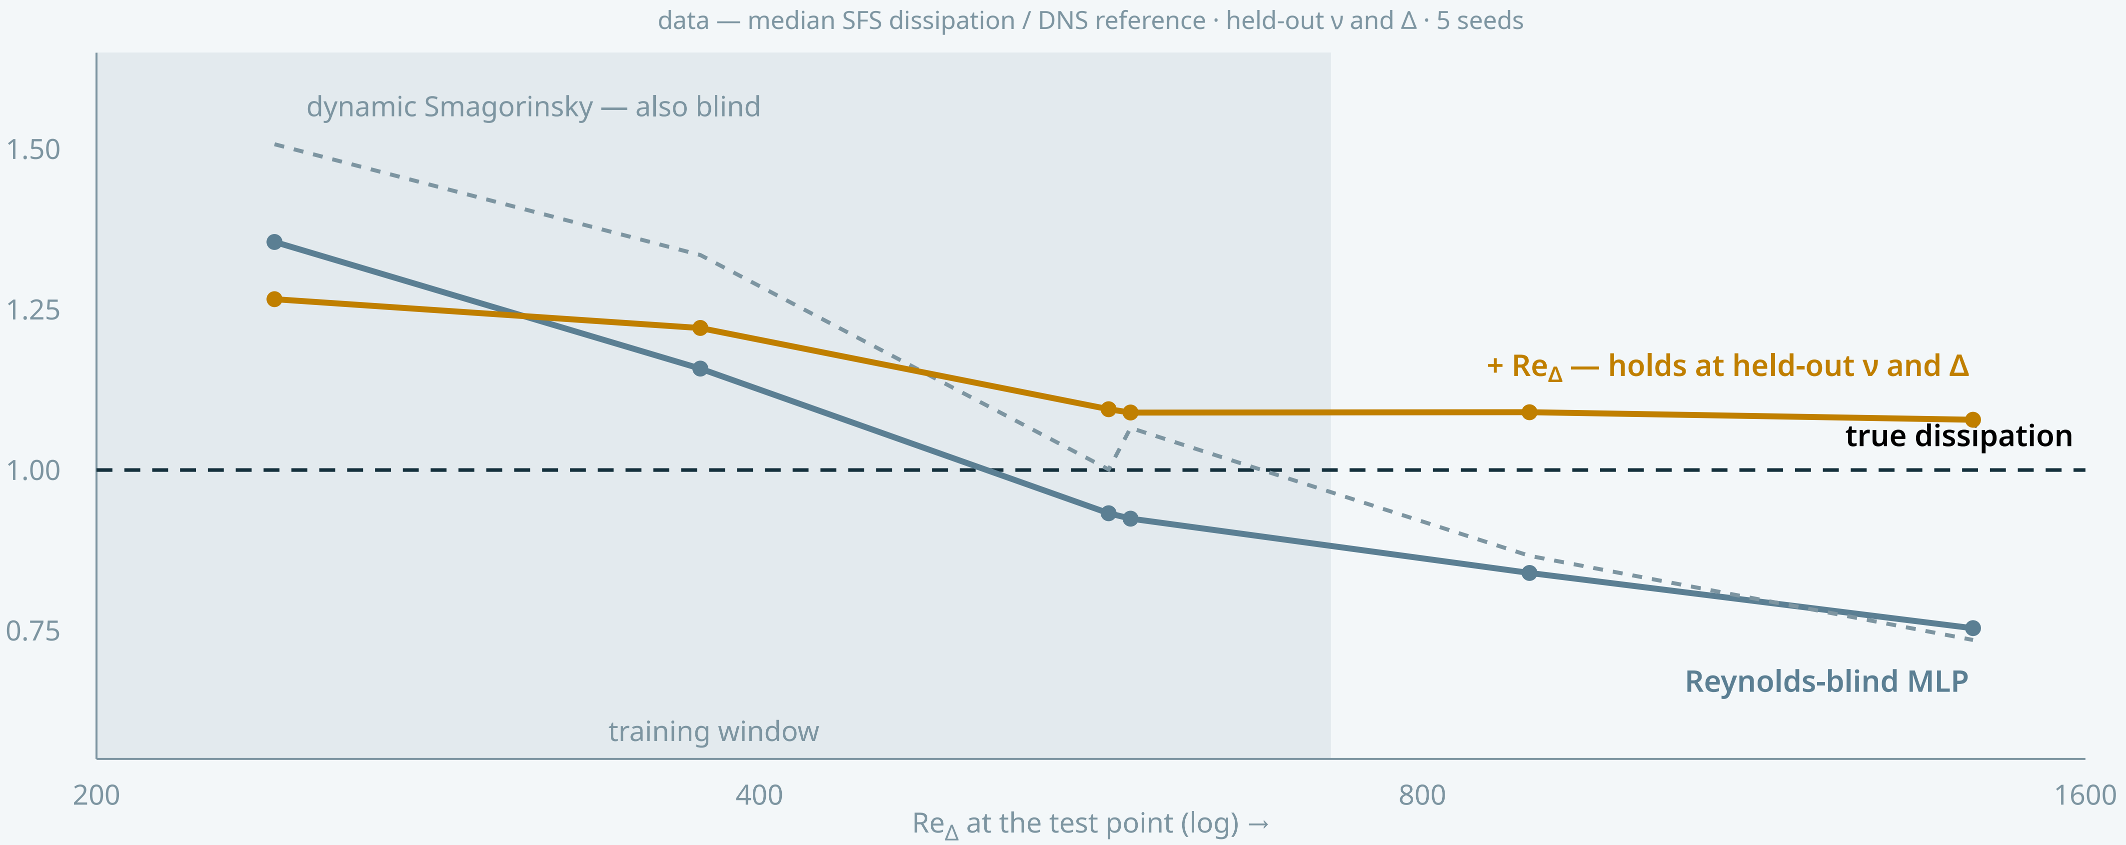

A flow it has never seen: Taylor–Green becoming turbulent

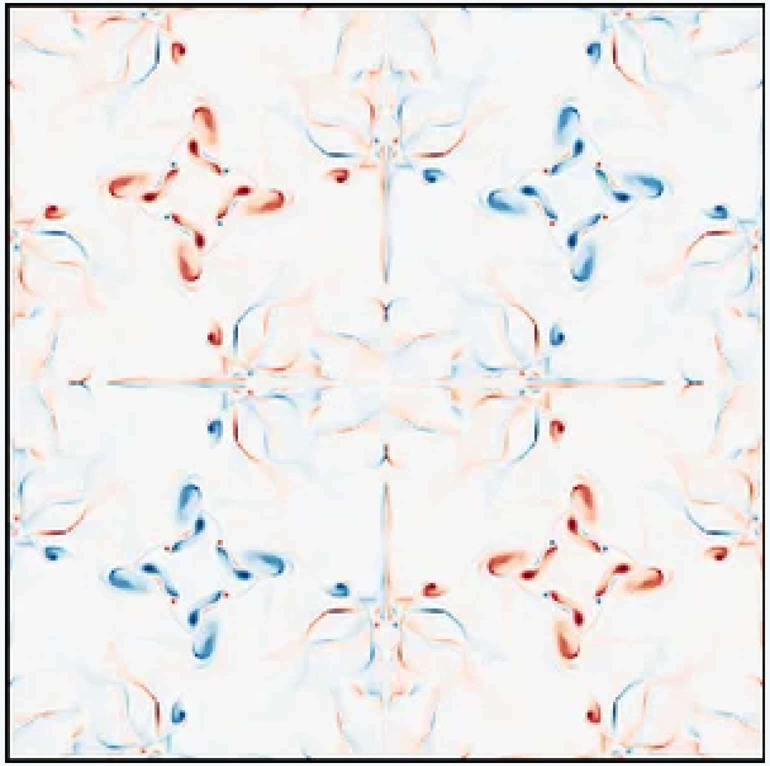

VORTICITY · $t = 0$

decaying, transitional — out of distribution twice over

laminar → transition → decay

press → to watch

all closures stay stable — but over-dissipate

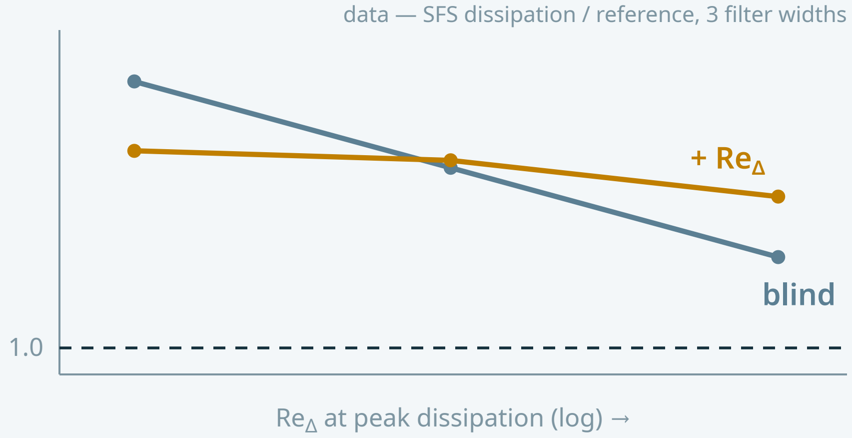

Re_Δ corrects the Reynolds-driven share of the excess, and honestly leaves the share owed to different flow physics

Couldn't a better architecture have fixed this? We checked.

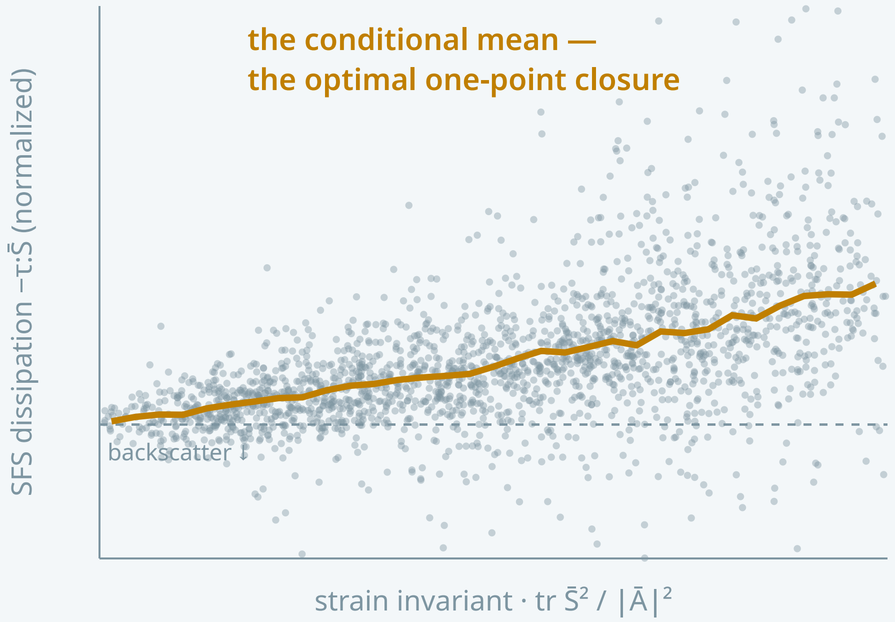

real data — one slice · one input value ↦ a cloud of true stresses

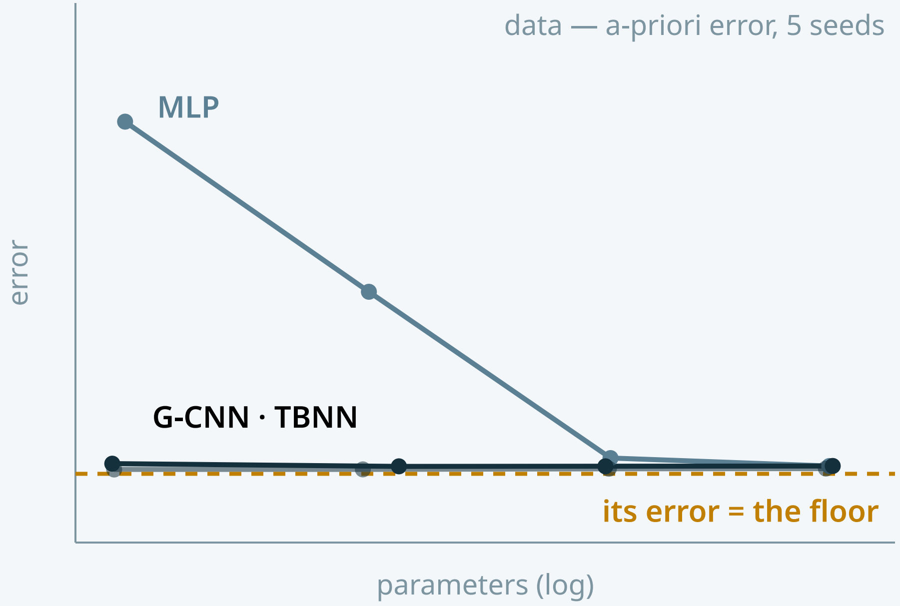

every architecture saturates to the **same** one-point optimum

so the Reynolds fix could never come from the architecture — **it had to come from the input**, available to all of them

TAKEAWAY

A closure that knows *what it's missing* works where it has never been.

Agdestein & Sanderse · CWI Amsterdam · "Approaching the optimal closure" · preprint & code available

the interactive story of this work
agdestein.github.io

CLOSE