Data-driven discrete closure models for large-eddy simulation of incompressible turbulence

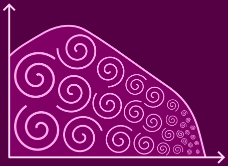

Syver Døving Agdestein

The problem: Simulation of turbulence

Turbulence: an ancient problem

250 BC

1500 AD

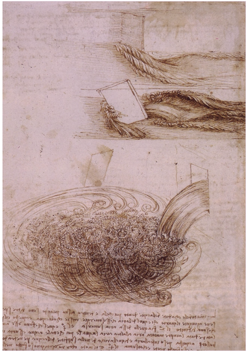

2000 AD

Boeing 777-200 photograph by Jules Meulemans,

licensed under the GNU Free Documentation License, version 1.2 or later.

Simulation of turbulence

The incompressible Navier-Stokes equations

Conservation of mass

$$\nabla \cdot u = 0$$

Conservation of momentum

$$\partial_t u + \nabla \cdot (u u - \nu \nabla u + p \delta) = f$$

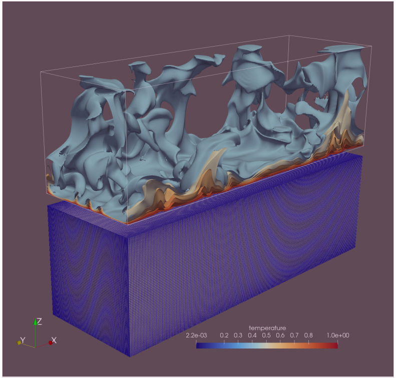

Lifecycle of turbulence

The Burgers' equation (1D)

$$\partial_t u = \underbrace{\nu \partial_{xx} u}_{\text{Diffusion}} - \underbrace{\partial_x \left(u^2/2\right)}_{\text{Convection}}$$

Effect on a Fourier mode

$$u = \mathrm{e}^{\mathrm{i}kx}$$

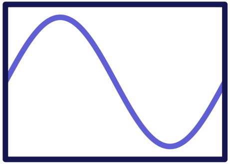

$$\partial_{xx} u = -k^2 \mathrm{e}^{\mathrm{i}kx}$$

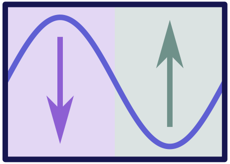

$$\partial_x \left(u^2/2\right) = \mathrm{i}k\mathrm{e}^{\mathrm{i}(2k)x}$$

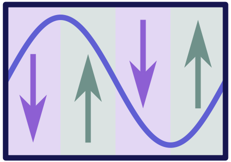

Spectral distribution of energy

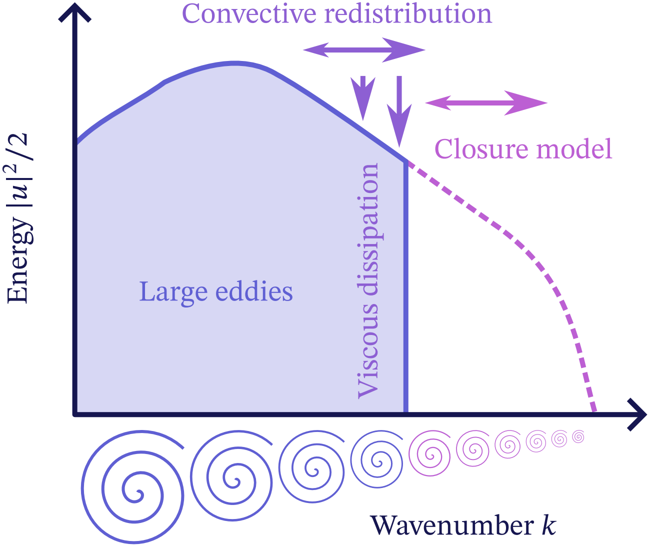

Large-eddy simulation (LES)

Filter out length scales smaller than Δ

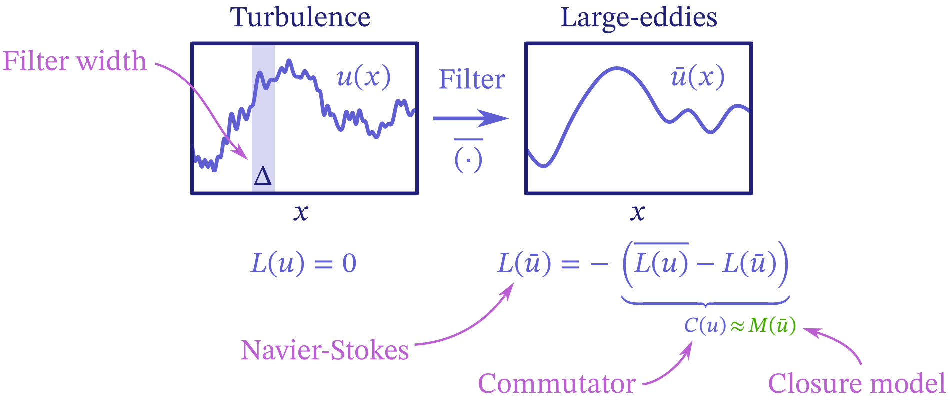

$$L(u) = 0$$

$$L(\bar{u}) = - \underbrace{\left(\overline{L(u)} - L(\bar{u}) \right)}_{C(u) \approx M(\bar{u})}$$

Issues with current closure models

- Closure model M_θ depends on parameters θ
- Learn θ from high-fidelity data $\mathcal{D}$:

$$\min_\theta \mathbb{E}_{u \sim \mathcal{D}} \left\| M_\theta(\bar{u}) - C(u) \right\|^2$$

- Models that fit target data well can still be **unstable** when deployed in LES

> *There are other commutators than $C(u)$.*
> *They cause the instabilities.*

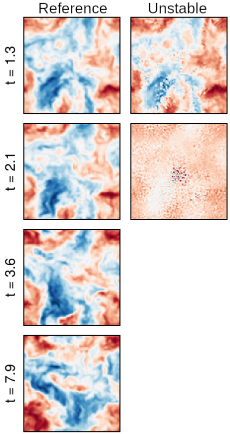

PhD research

Identify all commutators

- from the filter $\overline{(\cdot)}$
- from the discretization ∇_h
- from symmetry groups in the model M_θ

Address all commutators

- Eliminate them by choosing $\overline{(\cdot)}$, ∇_h, and M_θ in a smart way
- Account for the remaining ones by tuning the parameters θ

Discretization-consistent LES

1. The problem: Simulation of turbulence
2. Discretization-consistent LES
3. Are we modeling the wrong stress tensor?
4. Time integration as filtering
5. Inductive bias and Reynolds-number generalization
6. Conclusion

A new framework: discretize first

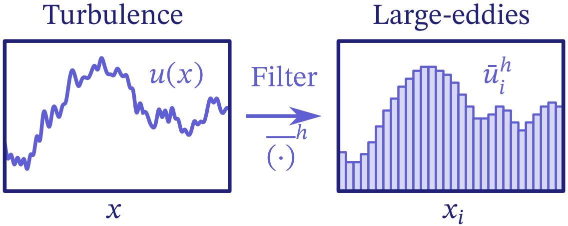

$$L(u) = 0$$

$$L_h(\bar{u}^h) = -\underbrace{\left(\overline{L(u)}^h - L_h(\bar{u}^h)\right)}_{C_h(u) \approx M_h(\bar{u}^h)}$$

Discretized Navier-Stokes

Commutator

Closure model

New filters, new stresses

Discrete filters

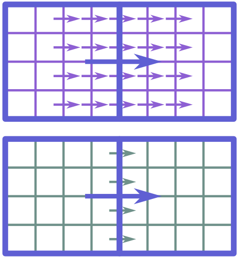

- Existing discrete filters induce discrete commutator errors.
- New discrete filters **restore** commutation (discrete integration-by-parts properties).

Stress tensors in LES

- Continuous LES stress expressions are symmetric and local in u.
- New discrete LES stress expressions are **non-symmetric** and **non-local** in u.

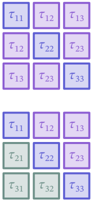

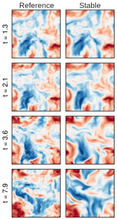

Are we modeling the wrong stress tensor?

1. The problem: Simulation of turbulence
2. Discretization-consistent LES
3. Are we modeling the wrong stress tensor?
4. Time integration as filtering
5. Inductive bias and Reynolds-number generalization
6. Conclusion

LES and FVM: the toolbox

Filtering

$$\overline{(\cdot)}^{\Delta}, \quad \overline{(\cdot)}^{h}$$

Closure modeling

$$\tau^{\Delta}(u) \approx m^{\Delta}(\bar{u}^{\Delta})$$

Pressure projection

$$\operatorname{div} \circ \pi = 0, \qquad \pi \pi = \pi$$

Divergence theorem

$$\partial_x^h = \overline{(\cdot)}^{h} \circ \partial_x$$

Flux reconstruction

$$\sigma(u) \approx \sigma^h(\bar{u}^h)$$

Time integration

$$u^* = u^n + \Delta t \dots$$
$$u^{n+1} = \pi u^*$$

Agdestein, Verstappen, Sanderse, *J. Comput. Phys.* 556 (2026) 114810

Filtering

LES filter

$$\bar{u}^{\Delta}(x) = \int_{\Omega} G^{\Delta}(x - y) u(y) \, \mathrm{d}y$$

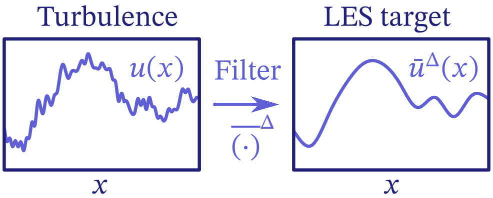

FVM filter

$$\bar{u}^{h}(x) = \frac{1}{h} \int_{x-h/2}^{x+h/2} u(y) \, \mathrm{d}y$$

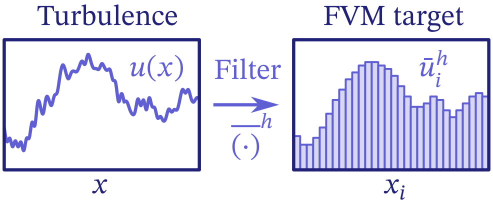

Divergence theorem

The mechanism powering the FVM

$$\partial_x^h u(x) = \frac{u\left(x + \frac{h}{2}\right) - u\left(x - \frac{h}{2}\right)}{h} = \frac{1}{h} \int_{x-h/2}^{x+h/2} \partial_y u(y) \, \mathrm{d}y = \overline{\partial_x u}^h (x)$$

The discrete divergence ∂_x^h is induced by the FVM filter $\overline{(\cdot)}^h$

$\implies$ exact discrete conservation laws for $\bar{u}^h$

Burgers' equation

$$\partial_t u + \partial_x \sigma(u) = 0, \quad \sigma(u) = uu/2 - \nu \partial_x u$$

Apply LES/FVM filters

$$\partial_t \bar{u}^{\Delta,h} + \partial_x^h \overline{\sigma(u)}^{\Delta} = 0$$

Introduce numerical flux

$$\partial_t \bar{u}^{\Delta,h} + \partial_x^h \sigma^h(\bar{u}^{\Delta,h}) = -\partial_x^h \tau^{\Delta,h}(u)$$

Unresolved discrete flux

$$\tau^{\Delta,h}(u) = \overline{\sigma(u)}^{\Delta} - \sigma^h(\bar{u}^{\Delta,h})$$

Discretization contribution

How much of the sub-filter flux comes from the discretization?

$$\tau^{\Delta,h}(u) = \tau^{\Delta,h}_{\text{classic}}(u) + \tau^{\Delta,h}_{\text{flux}}(u) + \tau^{\Delta,h}_{\text{div}}(u)$$

$$\begin{aligned}
\tau^{\Delta,h}_{\text{classic}}(u) &:= \overline{\sigma(u)}^{\Delta,h} - \sigma(\bar{u}^{\Delta,h}) \\
\tau^{\Delta,h}_{\text{flux}}(u) &:= \sigma(\bar{u}^{\Delta,h}) - \sigma^{h}(\bar{u}^{\Delta,h}) \\
\tau^{\Delta,h}_{\text{div}}(u) &:= \overline{\sigma(u)}^{\Delta} - \overline{\sigma(u)}^{\Delta,h}
\end{aligned}$$

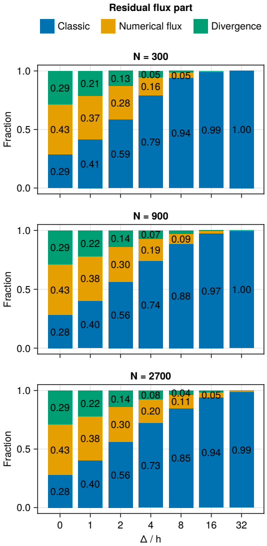

Discretization-aware closure

Fitting the Smagorinsky constant for Burgers' equation

$$m(u, \theta) = -\theta s(u), \quad m^h(u, \theta) = -\theta s^h(u),$$

$$s(u) = \left|\partial_x u\right| \partial_x u, \quad s^h(u) = \left|\partial_x^h u\right| \partial_x^h u,$$

$$\theta_{\text{classic}} = -\frac{s(\bar{u})\tau(u)}{s(\bar{u})s(\bar{u})},$$

$$\theta_{\text{informed}} = -\frac{s^h(\bar{u})\tau^h(u)}{s^h(\bar{u})s^h(\bar{u})}.$$

> *Comparison:*
> $m^h(\cdot, \theta_{\text{classic}})$ *vs* $m^h(\cdot, \theta_{\text{informed}})$

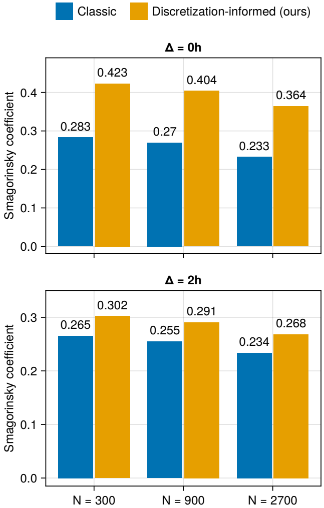

Discretization-aware closure

Errors and spectra

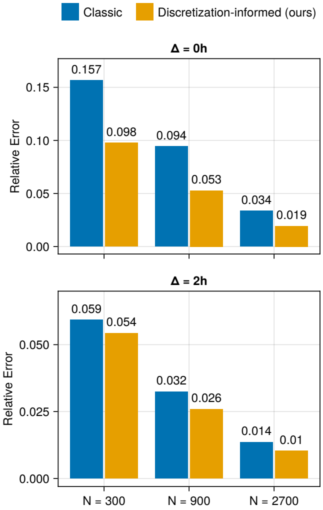

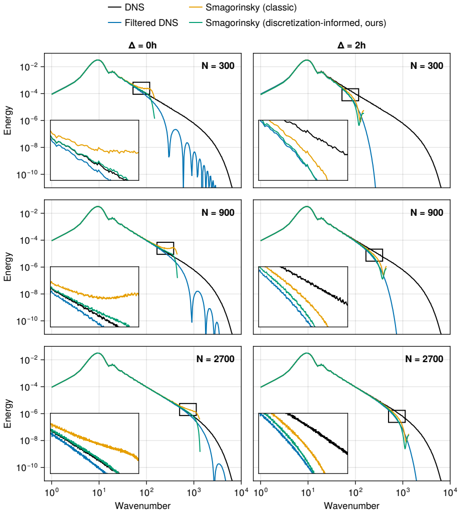

Incompressible Navier-Stokes

Projection form

$$\partial_t u = -\nabla \cdot \pi \sigma(u), \quad \sigma(u) = u u - \nu \left(\nabla u + \nabla u^T \right)$$

Exact LES equations

Infinitesimal LES: $\partial_t \bar{u} = -\overline{\nabla \cdot \pi \sigma(u)} = -\nabla \cdot \overline{\pi \sigma(u)}$

Discrete LES: $\partial_t \bar{u}^h = -\overline{\nabla \cdot \pi \sigma(u)}^h = -\nabla_h \cdot \overline{\pi \sigma(u)}^{h,*}$

> *The exact LES-FVM stress is non-symmetric and non-local in u.*

DNS-aided LES

Stress expressions for incompressible Navier-Stokes

LES equation

$$\partial_t \bar{u}^h = -\nabla_h \cdot \pi^h \left[\sigma^h(\bar{u}^h) + \tau^h(u)\right]$$

Five stress candidates

- τ^h_A: No SFS
- τ^h_B: Classical SFS
- τ^h_C: Classical SFS with num. flux
- τ^h_D: Correct SFS
- τ^h_E: Correct SFS (symmetrized)

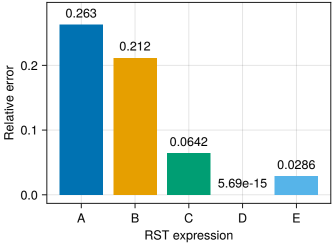

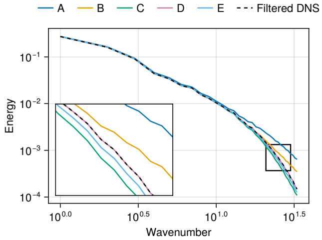

Take-away: the exact discrete stress

The discretization is a fundamental part of LES.

For incompressible Navier-Stokes, the LES-FVM stress tensor is non-local and non-symmetric.

Infinitesimal structural closure models have the wrong structure; infinitesimal functional closure models have the wrong dissipation profile.

Implicit and explicit LES unified into one coherent framework.

Time integration as filtering

Forward Euler is a filter

Filter-swap in space (before)

$$\partial_x^h = \overline{(\cdot)}^h \circ \partial_x$$

Filter-swap in time (new)

$$\bar{u}^\tau(t) := \frac{1}{\tau}\int_0^\tau u(t+s)\,\mathrm{d}s$$

$$\partial_t^\tau u(t) := \frac{u(t+\tau) - u(t)}{\tau} = \overline{\partial_t u}^\tau(t)$$

The temporal filter-swap

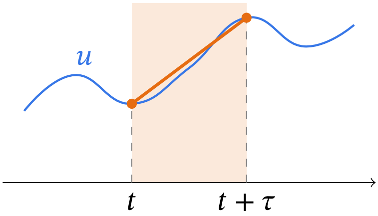

Chord slope $\partial_t^\tau u$ = mean of $\partial_t u$ over the step

Agdestein, *Time integration as filtering: a space-time discretization-aware LES formulation*, DLES15 (2026)

> *Forward Euler is a one-sided top-hat time filter.*

Space-time filter-swap

An exact fully discrete conservation law

Filter the conservation law in space *and* time

$$\partial_t u + \partial_x \sigma(u) = 0$$

$$\implies \quad \partial_t^\tau \bar{u}^h + \partial_x^h \overline{\sigma(u)}^\tau = 0$$

> *Exact*: *forward Euler driven by the step-averaged flux* $\overline{\sigma(u)}^\tau$ *reproduces* $\bar{u}^h$ *at every time node – at any CFL number.*

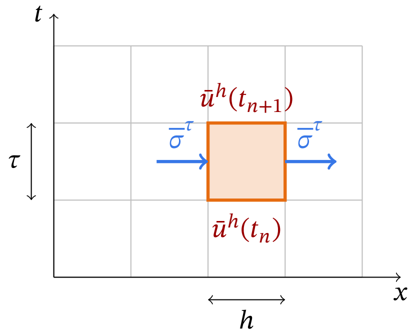

One space-time cell: update = time-averaged fluxes through the faces

The residual flux gains exactly one term

Space-time decomposition of the closure target

$$\tau^{h,\tau}(u) := \overline{\sigma(u)}^{\tau} - \sigma^h(\bar{u}^h) = \underbrace{\tau^h_{\text{classic}}(u) + \tau^h_{\text{flux}}(u) + \tau^h_{\text{div}}(u)}_{\text{spatial terms (as before)}} + \underbrace{\tau^\tau_{\text{time}}(u)}_{:= \overline{\sigma(u)}^{\tau} - \sigma(u)}$$

> *No temporal commutator appears: the new term is a pure flux time-quadrature error, shrinking with the order of the integrator.*

Leading order: Lax-Wendroff diffusion

$$\tau^\tau_{\text{time}}(u) = \frac{\tau}{2}\,\partial_t \sigma(u) + \mathcal{O}(\tau^2) = -\frac{\tau}{2}\,\left(\sigma'(u)\right)^2 \partial_x u + \mathcal{O}(\tau^2)$$

Eddy viscosity $\frac{\tau}{2}\left(\sigma'(u)\right)^2 \geq 0$: dissipative by construction – the diffusion that stabilizes forward-Euler convection

A space-time-aware closure

Augment the discretization-aware Smagorinsky model

$$m^{h,\tau}(\cdot) := \underbrace{-\theta_h^2 h^2 \left|\partial_x^h \cdot\right| \partial_x^h \cdot}_{\text{spatial terms}} \underbrace{-\,\theta_\tau\, \tau \left(\sigma'(\cdot)\right)^2 \partial_x^h \cdot}_{\text{temporal term}}$$

- One new coefficient θ_τ, with theoretical value $\frac{1}{2}$ from the Lax-Wendroff expansion
- Fitted a priori by dissipation matching: the measured temporal dissipation is linear in τ, as predicted

A priori: the temporal term dominates at practical CFL

Burgers turbulence, forward Euler, CFL 0.06 – 0.89

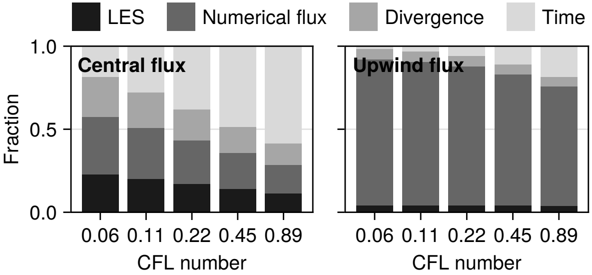

- Spatial terms: fixed proportions, independent of the time step
- Temporal share grows with CFL: 0.19 → 0.59 (central flux)
- At practical CFL numbers, the **largest single term** is the one space-only formulations discard

A posteriori: a CFL-robust closure

Coarse forward-Euler LES, error at final time

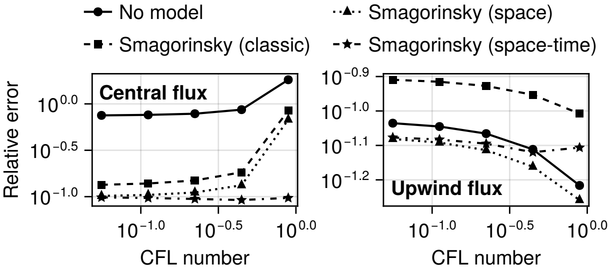

- Step-agnostic closures degrade as the CFL number grows
- The space-time-aware closure: error independent of CFL (0.092 – 0.098, central flux)
- Upwind flux over-dissipates; forward Euler partially cancels it

Inductive bias and Reynolds-number generalization

Symmetries of the Navier-Stokes equations

Symmetries of the Navier-Stokes equations

Galilean invariance

Symmetries of the Navier-Stokes equations

Rotational invariance

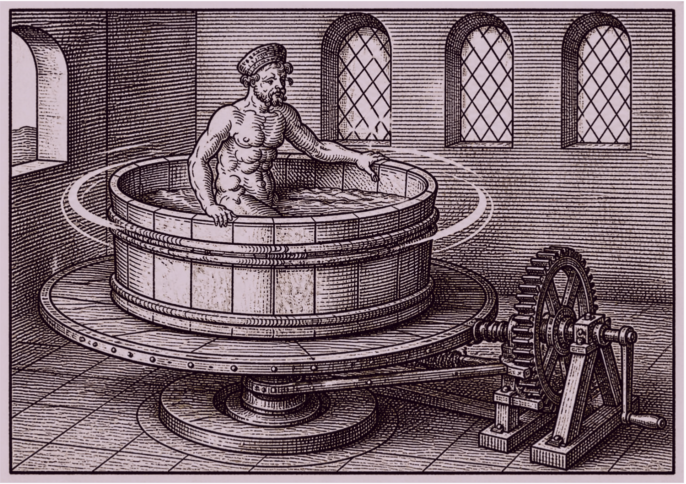

Symmetries of the Navier-Stokes equations

Reflectional invariance

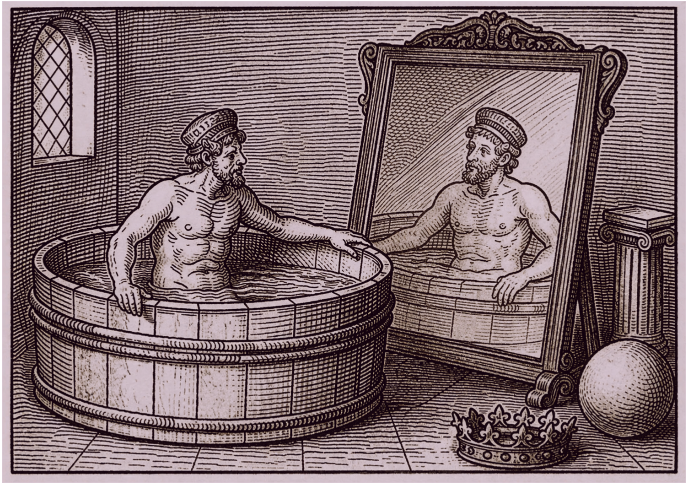

Symmetries of the Navier-Stokes equations

Scaling invariance

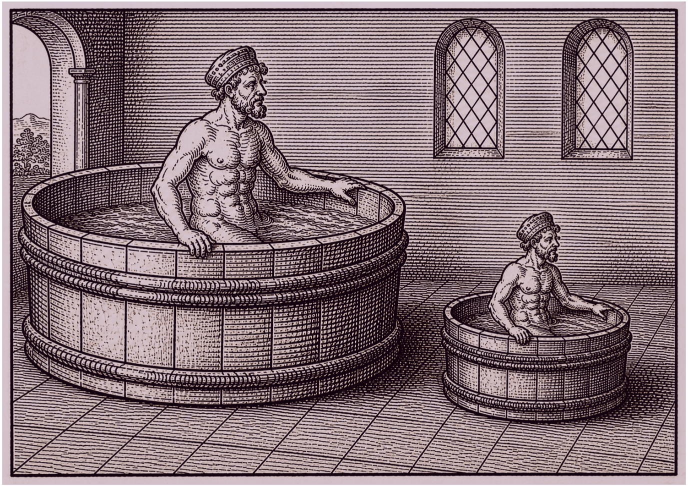

Symmetries of the Navier-Stokes equations


Galilean

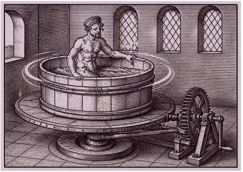
Rotational


Reflectional

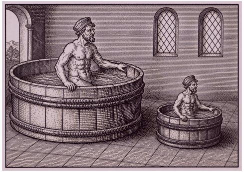
Scaling

> *G is a symmetry-group of the Navier-Stokes equations L if and only if*
>
> $$L(gu) = gL(u)$$
>
> *for all velocities u and group elements $g \in G$.*

> *How to choose M_θ such that $M_\theta(g\bar{u}) = gM_\theta(\bar{u})$ for all θ, $\bar{u}$, and $g \in G$?*

Which symmetries survive discretization?

Symmetry	Navier-Stokes	Discrete LES + pointwise closure
Time translation	preserved	preserved
Space translation	preserved	preserved (pointwise closure)
Constant Galilean	preserved	preserved (input $\nabla \bar{u}$, not $\bar{u}$)
Generalized Galilean	preserved	broken (frame-fixed forcing)
Rotation/reflection $O(3)$	preserved	reduced to octahedral G, $\lvert G \rvert = 48$
Scaling (two-parameter)	preserved	reduced to $a^2 = b$ (fixed Δ, hidden ν)
Pressure gauge	preserved	preserved (deviatoric closure)

> *A closure can only meaningfully preserve the symmetries that the discrete equations retain: on a grid, the octahedral group is the ceiling.*

Agdestein, Sanderse, *Approaching the optimal closure: inductive bias and Reynolds-number generalization in data-driven LES* (2026)

Effect of symmetry-preserving models

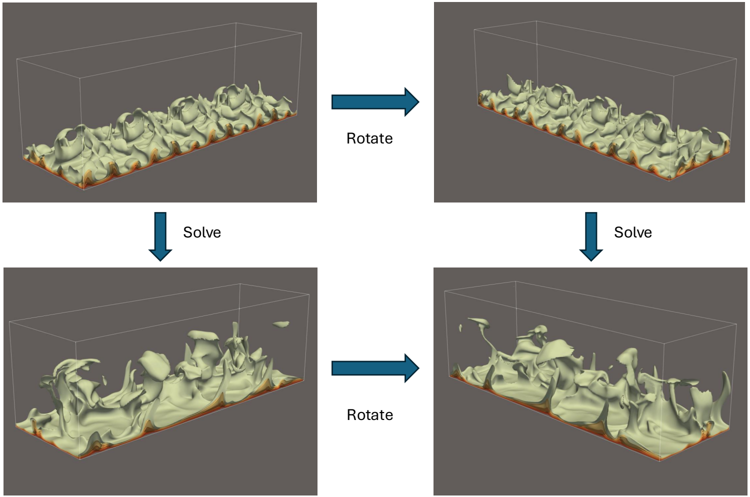

Rotate

Solve

Solve

Rotate

Three closures, one construction

Same inputs, same outputs, parameter-matched: only the inductive bias differs

Shared pointwise construction

$$M_\theta(\bar{u}) = \underbrace{\Delta^2 \, \|\nabla \bar{u}\|^2}_{\text{scaling}} \; \mathrm{NN}_\theta\Big(\underbrace{\nabla \bar{u} \,/\, \|\nabla \bar{u}\|}_{\text{Galilean}}\Big)$$

MLP

Unconstrained baseline: no roto-reflection equivariance

G-CNN

Equivariance through the ***architecture****: weights tied over the 48 octahedral group elements*

TBNN

Equivariance through the ***features****: invariants* λ_k *and tensor basis* T_k

Tensor basis neural network

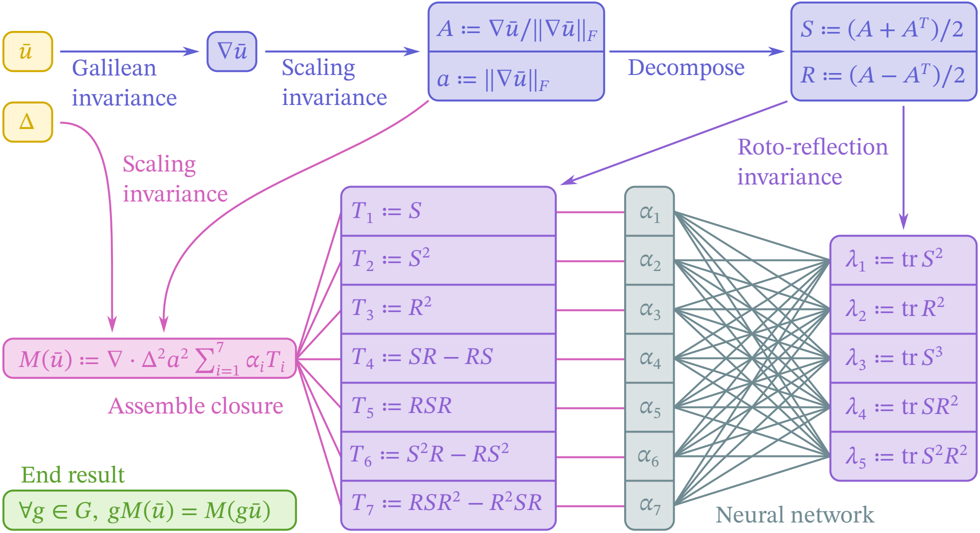

$\bar{u}$ → Galilean invariance → $\nabla \bar{u}$ → Scaling invariance →

$A := \nabla \bar{u} / \|\nabla \bar{u}\|_F$

$a := \|\nabla \bar{u}\|_F$

→ Decompose →

$S := (A + A^T)/2$

$R := (A - A^T)/2$

Δ

Scaling invariance

Roto-reflection invariance

$T_1 := S$

$T_2 := S^2$

$T_3 := R^2$

$T_4 := SR - RS$

$T_5 := RSR$

$T_6 := S^2R - RS^2$

$T_7 := RSR^2 - R^2SR$

α_1, α_2, α_3, α_4, α_5, α_6, α_7

$\lambda_1 := \operatorname{tr} S^2$

$\lambda_2 := \operatorname{tr} R^2$

$\lambda_3 := \operatorname{tr} S^3$

$\lambda_4 := \operatorname{tr} SR^2$

$\lambda_5 := \operatorname{tr} S^2R^2$

Neural network

$M(\bar{u}) := \nabla \cdot \Delta^2 a^2 \sum_{i=1}^{7} \alpha_i T_i$

Assemble closure

End result

$\forall g \in G, \; gM(\bar{u}) = M(g\bar{u})$

Group-convolutional weight sharing

- **Synthesize** equivariant weights directly (gather), instead of **projecting** (group-averaging): exactly equivariant in floating-point arithmetic
- Same weight sharing as a classical CNN, but over roto-reflections instead of shifts – no locality, so the blocks are dense

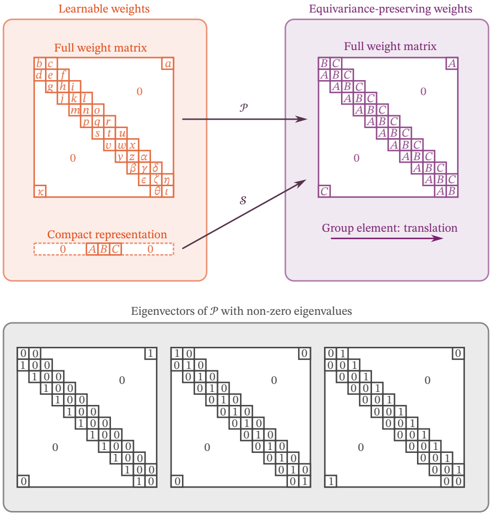

A posteriori: forced isotropic turbulence

Energy and sub-filter dissipation over the evaluation window

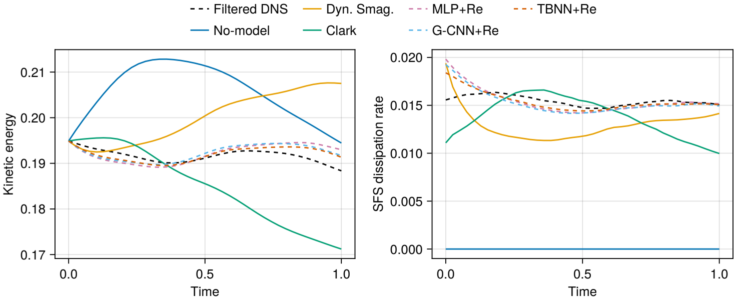

- The three data-driven closures track the reference most closely – and are indistinguishable from each other
- No-model accumulates energy; Clark is marginally stable; dynamic Smagorinsky drifts

A posteriori: spectra and backscatter

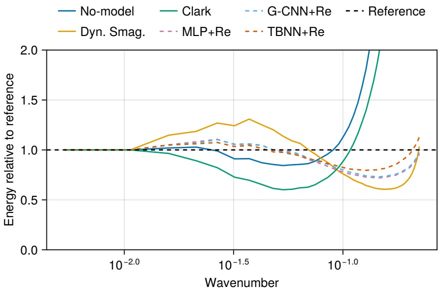

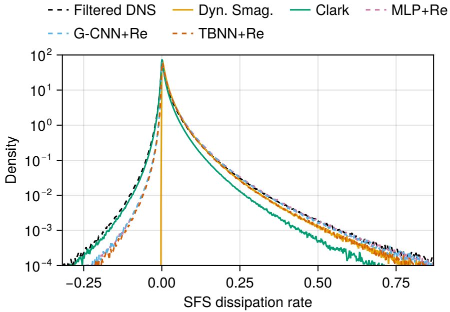

- Dynamic Smagorinsky: no backscatter by construction; Clark: backscatter but too little forward transfer
- Data-driven closures reproduce both tails of the sub-filter dissipation

Inductive bias buys parameter efficiency, not accuracy

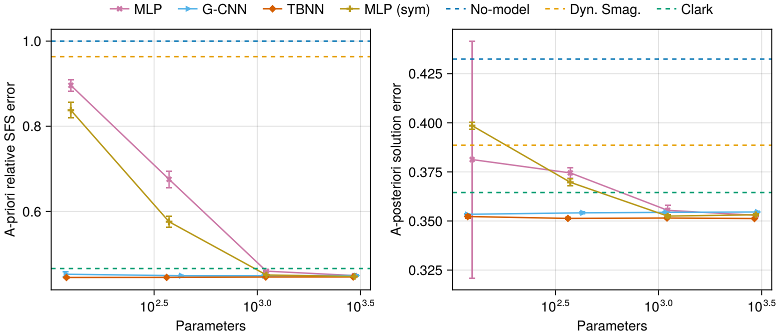

- TBNN and G-CNN are already saturated at ~ 120 parameters; the MLP needs $\sim 10\times$ more to catch up
- All three converge to the same floor: the one-point *optimal closure* $\mathbb{E}\left[\tau(u) \mid \nabla \bar{u}\right]$ (Langford & Moser 1999)

Generalization across Reynolds numbers

Train at three viscosities, test on held-out ν and Δ/h

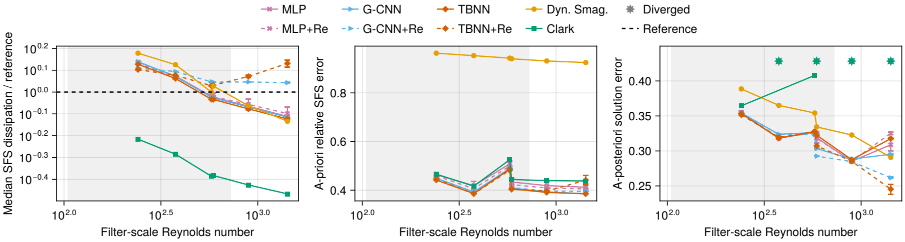

- The stress *structure* is Reynolds-independent (flat a priori error), but the *dissipation calibration* drifts outside the training band
- Feeding one scalar, $\mathrm{Re}_\Delta = \Delta^2 \|\nabla \bar{u}\| / \nu$, restores the calibration – for all three architectures alike

Out-of-distribution: decaying Taylor-Green vortex

Trained closures applied unchanged to an unforced, transitional flow

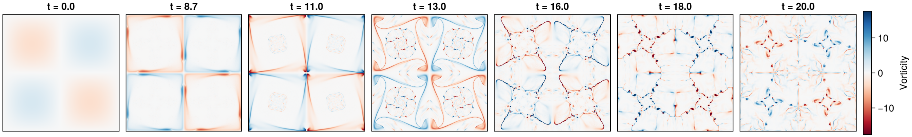

- Laminar roll-up → transition → viscous decay: none of it appears in the forced training data
- Reuse is legitimate: the closures are amplitude- and Reynolds-invariant *in form*

Taylor-Green: partial generalization

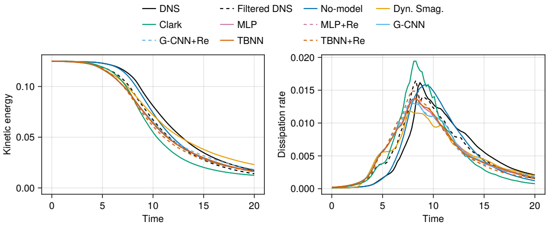

- Data-driven closures stay stable and structurally faithful through the transition
- But they carry a flow-type over-dissipation ($\sim 2\times$) in this milder regime
- Re_Δ corrects the Reynolds share of the drift, not the flow-type share

Take-away: does equivariance matter?

> *Equivariance costs no accuracy: even the unconstrained MLP learns a near-equivariant map from isotropic data.*

> *The inductive bias buys parameter efficiency and reproducibility, not a lower error floor: all pointwise closures saturate at the same optimal closure.*

> *Reynolds-number generalization comes from the* Re_Δ *input, not from the architecture.*

Conclusion

Take-away

> *The main problem in LES is the presence of commutator errors.*

> *We identify new commutators in LES arising from the filter, discretization, and symmetry-groups.*

> *We cancel out commutators through new filters, discrete formulations, and model architectures.*

> *Exact expressions for the remaining commutators allow for producing unbiased training data for data-driven closure models.*

This thesis offers a mathematically principled foundation for reducing the computational cost of turbulence simulations.

Postlude

PhD and PostDoc positions in *probabilistic turbulence*

`https://www.cwi.nl/en/jobs/vacancies/`

Thank you!

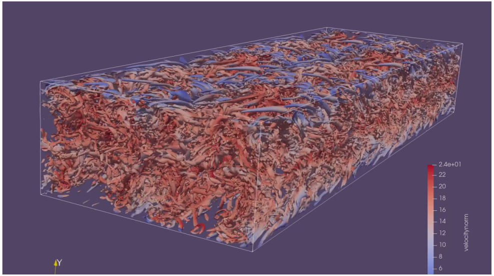
2.4e+01
22
20
18
16
14
12
10
8
6
velocitynorm
Y

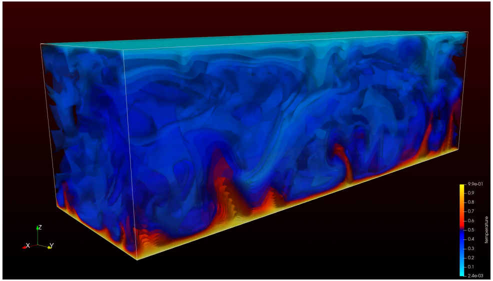
9.9e-01
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
2.4e-03
temperature
Z
X
Y

Data-driven discrete closure models for large-eddy simulation of incompressible turbulence

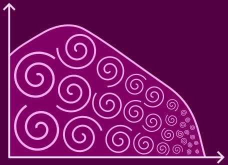

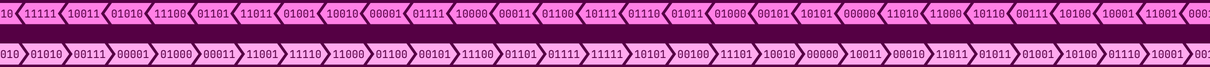

Syver Døving Agdestein