Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A generalised Spalart-Allmaras model with blended near-wall and free-shear coefficients: multi-objective calibration in an open-source framework

Tony Di Fabbio
tony.difabbio@polito.it
Politecnico di Torino
Riccardo Vallino
Politecnico di Torino
Eric Segalerba
University of Salerno
Viola Rossano
Leonardo (Italy)
Joel Guerrero
Leonardo (Italy)
Andrea Ferrero
Politecnico di Torino

Research Article

Keywords: Spalart-Allmaras, Reynolds-averaged turbulence modelling, coefficient blending, multi-objective calibration, data-driven turbulence modelling

Posted Date: October 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11333768/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A generalised Spalart-Allmaras model with blended near-wall and free-shear coefficients: multi-objective calibration in an open-source framework

Tony Di Fabbio[1*], Riccardo Vallino[1], Eric Segalerba[2], Viola Rossano[3], Joel Guerrero[4], Andrea Ferrero[1]

[1*]Department of Mechanical and Aerospace Engineering (DIMEAS), Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, 10129, Italy.
[2]Department of Industrial Engineering (DIIN), University of Salerno, Via Giovanni Paolo II 132, Fisciano (SA), 84084, Italy.
[3]Digital Twin Laboratory, Digital Engineering, Leonardo S.p.A., C.so Castelfidardo 22, Torino, 10138, Italy.
[4]Digital Twin Laboratory, Digital Engineering, Leonardo S.p.A., Via Pieragostini 80, Genova, 16151, Italy.

*Corresponding author(s). E-mail(s): tony.difabbio@polito.it;

Abstract

The Spalart-Allmaras model owes its robustness and wide industrial use to a few globally constant coefficients, calibrated on equilibrium wall-bounded flows; the same uniformity limits its accuracy in separated and free-shear regions. This work explores how far the Spalart-Allmaras model can be generalised without losing that robustness. Each calibration coefficient receives a near-wall and a free-shear value, blended locally by a shielding function from hybrid Reynolds-averaged/large-eddy methods, so that one model adapts to the local flow regime without zonal switching. Both coefficient states are calibrated by multi-objective optimisation in an open-source solver, on two benchmark suites spanning attached, free-shear, separated and corner flows. Attached flows prove insensitive to the near-wall state, so the generalised model preserves their accuracy while improving where the baseline model struggles, reducing the jet objective twentyfold and the separated-hump objective by a fifth. The flows disagree mainly on the free-shear production, which sets the transferability of each calibration. The approach also separates what coefficients can fix from what is

structural: no coefficient set generates the anisotropy-driven secondary flow of a square duct, and a topology-sensitive shielding only relieves the associated trade-off, whereas a calibrated quadratic constitutive relation reduces the duct objective and produces the corner vortices. On three unseen configurations the model calibrated on external flows retains the accuracy of the baseline model.

Keywords: Spalart-Allmaras, Reynolds-averaged turbulence modelling, coefficient blending, multi-objective calibration, data-driven turbulence modelling

1 Introduction

Reynolds-averaged Navier-Stokes (RANS) turbulence models remain the dominant tool of industrial computational fluid dynamics (CFD), as a practical compromise between cost, robustness and accuracy. The one-equation Spalart-Allmaras (SA) model (Spalart and Allmaras, 1992) is among the most widely used in external aerodynamics, turbomachinery and general engineering: it solves a single transport equation for a modified eddy viscosity, at modest cost, with a calibration that is reliable across a broad range of attached wall-bounded flows. The implementation clarifications and "negative" formulation of Allmaras et al (2012) further improved its stability and portability. Its coefficients are globally constant and were selected to represent both equilibrium boundary layers and free-shear flows. This gives a portability rarely matched by more complex closures, but a single set is necessarily a compromise between regimes with markedly different turbulence physics, and accuracy deteriorates where the flow departs from the equilibrium assumptions of the calibration. Such limitations are well documented for strong adverse-pressure-gradient separation, large-scale detachment, streamline curvature, rotation and turbulence-driven secondary motion (Demuren and Rodi, 1984; Spalart, 2000).

Efforts to improve eddy-viscosity models have followed two complementary strategies. The first adds targeted corrections and keeps the baseline calibration. Within the SA family, examples are the rotation/curvature correction of Shur et al (2000), which modifies production through invariants of the strain-rate and rotation-rate tensors, and the quadratic constitutive relation (QCR) of Spalart (2000), hereafter QCR2000, which adds nonlinear tensor terms to the Boussinesq stress. Such corrections improve selected flow classes, but the calibration coefficients of the parent model remain fixed and globally uniform. The second strategy widens the calibration itself. The generalised k-ω model of Menter and Matyushenko (2025) recasts selected closure coefficients as tuning parameters tied to specific mechanisms (separation, mixing, near-wall behaviour), expanding the calibration space without altering the structure of the parent closure. Such models suggest that some limitations of eddy-viscosity closures stem less from the governing equations than from the assumption that one globally constant coefficient set suits regions of different turbulence physics.

Regime-dependent coefficients require a mechanism that identifies the local flow state. Hybrid RANS/large-eddy simulation (LES) methods have long addressed this problem: delayed detached-eddy simulation (DDES) (Spalart et al, 2006) introduced

a wall-protection (shielding) function that prevents premature RANS-to-LES transitions on refined grids, and improved DDES (IDDES) (Shur et al, 2008) added length-scale and grid-dependent criteria. We repurpose these shielding concepts to govern the transition between two coefficient states within a purely RANS model. Blending two coefficient sets is itself not new: the shear-stress transport (SST) model switches between the k-ω and k-ε constants through its wall-distance-based F_1 function (Menter, 1994), and the strain-adaptive linear SA (SALSA) formulation makes the SA production coefficient a function of the local strain (Rung et al, 2003). The novelty of this work is the combination: every SA calibration coefficient carries two states, blended by a DDES/IDDES-type shielding function within a one-equation RANS model and calibrated jointly by multi-objective optimisation. Coefficient adaptation cannot remove deficiencies that are structural to the closure. Turbulence-driven secondary motion in non-circular ducts is the classic example: it arises from Reynolds-stress anisotropy, which no linear eddy-viscosity closure reproduces, whatever its coefficients (Demuren and Rodi, 1984). QCR2000 was introduced for this problem, and Tamaki and Kawai (2024) revisited it with wall-resolved LES data of corner-flow separation, proposing the QCR2024 extension with further nonlinear stress terms and a recalibrated coefficient set.

High-fidelity numerical and experimental data have driven a broad effort in data-driven turbulence modelling, reviewed by Duraisamy et al (2019), and in uncertainty quantification, reviewed by Xiao and Cinnella (2019). The closest relatives of the present work treat the closure coefficients themselves as uncertain or tunable: Bayesian calibration of the k-ε coefficients across flow scenarios (Edeling et al, 2014b,a), the epistemic uncertainty carried by the free coefficients of RANS models (Margheri et al, 2014), and the propagation of SA closure-coefficient intervals through transonic wall-bounded flows (Schaefer et al, 2017). These works showed that coefficient sensitivities are strongly flow-dependent, which underpins regime-dependent calibration. A second family infers spatial correction fields and regresses them on local flow features: field inversion and machine learning (Parish and Duraisamy, 2016), applied to SA on separated aerofoil flows by Singh et al (2017), and tensor-basis neural networks for the anisotropy (Ling et al, 2016). A third family discovers algebraic corrections symbolically: gene-expression programming of the stress-strain relation (Weatheritt and Sandberg, 2016), its CFD-driven variant with the RANS solver inside the training loop (Zhao et al, 2020; Di Fabbio et al, 2024), extended to multi-objective (Waschkowski et al, 2022; Kotzlowski et al, 2026) and multi-case (Fang et al, 2023) training, and the sparse-regression route to explicit algebraic Reynolds-stress models, SpaRTA (Schmelzer et al, 2020), later made probabilistic by sparse Bayesian learning (Cherroud et al, 2022). Shared training and assessment data now range from the curated RANS and direct numerical simulation (DNS) dataset of McConkey et al (2021) to the NASA data-driven turbulence-modelling challenge (Rumsey et al, 2022) and its successor organised by the European Research Community on Flow, Turbulence and Combustion (ERCOFTAC) (Buchanan et al, 2026).

The present work is data-driven in the same sense, with CFD inside the optimisation loop, but keeps the lowest-dimensional target: the calibration coefficients of the

parent model, split into two regime states. The question is how far SA can be generalised without losing the robustness that made it an industrial standard. The aim is not the best possible fit but to locate the deficiencies of the model, to identify the coefficients that drive each gain and to measure how much can be gained without complicating the formulation, so that interpretability, portability and implementation simplicity are preserved. In the resulting generalised SA model (GSA), each calibration coefficient has an Inner set for attached near-wall turbulence and an Outer set for detached and free-shear regions, blended continuously in space by the shielding function, without zonal decomposition or user-defined switching criteria. Both sets are calibrated by multi-objective evolutionary optimisation (Deb et al, 2002), seeded by particle-swarm optimisation (PSO) (Kennedy and Eberhart, 1995), on benchmark suites from the NASA Turbulence Modeling Resource (TMR) (Rumsey et al, 2010) and the ERCOFTAC validation databases (McConkey et al, 2021). Every objective is normalised by the error of baseline SA, so that values below unity are improvements and any loss of an existing capability is equally visible. Two independent calibrations, one per four-case suite, deliver GSA-N (NASA suite) and GSA-E (ERCOFTAC suite), with a QCR2024-augmented variant, GSA-E+QCR2024. Each model is assessed on the other suite, on three configurations that no calibration saw (an aerofoil polar through stall, a transonic shock-induced separation and a three-dimensional wing-body junction), and through a realisability analysis of the modelled Reynolds stress. On the ERCOFTAC suite, a conflict between separated-flow and square-duct predictions reflects not imperfect calibration but the inability of a linear eddy-viscosity closure to represent anisotropy-driven secondary motion. It motivates two developments: a topology-sensitive floor in the shielding function, to strengthen the separation response in rotation-dominated regions, and the calibration of the QCR2024 relation (Tamaki and Kawai, 2024). The work therefore separates, systematically, the limitations of coefficient selection from those of closure structure.

Section 2 introduces the generalised formulation and Section 3 the calibration framework, benchmark suites and campaigns. Section 4 presents the results in the order of the calibration strategy, and Section 5 concludes.

2 Generalised Spalart-Allmaras formulation

Attached boundary layers, separated flows and free-shear regions differ in their balance of turbulence production, destruction and diffusion, yet baseline SA applies one coefficient set throughout the domain. We therefore let selected calibration coefficients vary continuously between a near-wall and a free-shear state under a hybrid RANS/LES shielding function; the SA transport equation and its numerics are retained, and only the calibration space is extended.

The parent model solves a single transport equation for the modified eddy viscosity $\tilde{\nu}$ (Spalart and Allmaras, 1992; Allmaras et al, 2012),

$$\frac{D\tilde{\nu}}{Dt} = C_{b1}\tilde{S}\tilde{\nu} + \nabla\cdot\left(\frac{\nu+\tilde{\nu}}{\sigma_{\tilde{\nu}}}\nabla\tilde{\nu}\right) + \frac{C_{b2}}{\sigma_{\tilde{\nu}}}\left|\nabla\tilde{\nu}\right|^2 - C_{w1}f_w\left(\frac{\tilde{\nu}}{d}\right)^2, \tag{1}$$

with the standard auxiliary relations for $\tilde{S}$, f_w and the eddy-viscosity closure $\nu_t = \tilde{\nu} f_{v1}$ unchanged from Spalart and Allmaras (1992). The coefficients C_{b1}, $\sigma_{\tilde{\nu}}$, C_{b2} and C_{w1} that weight production, diffusion and destruction are those generalised below (C_{w1} through the reconstruction of Eq. (3)); the near-wall damping functions keep their original constants. The implementation follows the OpenFOAM SA model: the trip term and f_{t2} are omitted, $\tilde{S}$ is clipped at $0.3\,\Omega$, and $\tilde{\nu}$ is bounded at zero (no negative-$\tilde{\nu}$ branch). Three combinations have a physical reading used throughout: C_{b1} sets the rate at which the mean shear produces the eddy viscosity; the diffusion term can be rewritten as $(1/\sigma_{\tilde{\nu}})(\nu + \tilde{\nu})\nabla^2\tilde{\nu} + [(1 + C_{b2})/\sigma_{\tilde{\nu}}]\,|\nabla\tilde{\nu}|^2$ (for constant ν and uniform $\sigma_{\tilde{\nu}}$; with a blended $\sigma_{\tilde{\nu}}$, keeping $1/\sigma_{\tilde{\nu}}$ inside the divergence as in Eq. (1) adds $(\nu+\tilde{\nu})\nabla(1/\sigma_{\tilde{\nu}})\cdot\nabla\tilde{\nu})$, so that the front-spreading coefficient $D = (1+C_{b2})/\sigma_{\tilde{\nu}}$ governs how fast an eddy-viscosity front spreads into non-turbulent fluid, the mechanism that sets the growth rate of shear layers in the original calibration (Spalart and Allmaras, 1992). Finally, C_{w1} sets the near-wall destruction that balances production in the logarithmic layer. Because destruction acts only within a few wall distances d of a solid surface, a free shear layer is sensitive to the production and diffusion coefficients only, a wall-bounded one also to destruction.

2.1 Coefficient blending and shielding

Each blended coefficient ψ has an Inner value ψ_I and an Outer value ψ_O, combined locally as

$$\psi(\mathbf{x}) = \mathcal{F}_{SA}\,(\psi_I - \psi_O) + \psi_O, \tag{2}$$

where $\mathcal{F}_{SA} \in [0, 1]$ is the shielding function: $\mathcal{F}_{SA} = 1$ selects the Inner set and $\mathcal{F}_{SA} = 0$ the Outer set. The names refer to the two limits of a boundary layer, but the Inner state is not confined to the wall region: the shielding also returns $\mathcal{F}_{SA} \approx 1$ in the irrotational freestream and in the bulk of internal flows (in the square duct, $\mathcal{F}_{SA} > 0.5$ in 99.7% of the cells), so the Outer state is selected only where the flow is detached or free-shear. The blended coefficients are $\sigma_{\tilde{\nu}}$, C_{b1} and C_{b2}. The wall-destruction coefficient is reconstructed locally from their blended values (it equals the per-state values reported in Section 4.1 only where $\mathcal{F}_{SA} = 0$ or 1),

$$C_{w1} = \frac{C_{b1}}{\kappa^2} + \frac{1 + C_{b2}}{\sigma_{\tilde{\nu}}}, \tag{3}$$

which preserves the classical logarithmic-layer consistency relation in both limiting states, so that, whatever Inner set the optimiser chooses, the log-law balance of the attached boundary layer holds by construction. The Inner coefficients act on the departure from equilibrium (the boundary-layer wake region and its response to pressure gradients), not on the log law itself.

The shielding function combines two branches inspired by DDES and IDDES (Spalart et al, 2006; Shur et al, 2008),

$$\mathcal{F}_{SA} = \min\left(f_{d\beta}, f_{d\Omega}\right). \tag{4}$$

The first branch protects the attached boundary layer through a wall-distance-based sensor,

$$r_d = \frac{\nu_{\text{eff}}}{\kappa^2 d^2 |\nabla U|}, \qquad f_d = 1 - \tanh\left((c_d r_d)^3\right), \tag{5}$$

$$f_{\text{step}} = \min\left(2e^{-9\alpha^2}, 1\right), \qquad \alpha = \max\left(0.25 - \frac{d}{\Delta}, -5\right), \tag{6}$$

$$f_{d\beta} = \max\left(1 - f_d, f_{\text{step}}\right). \tag{7}$$

Here d is the wall distance, $|\nabla U| = (\partial_j U_i \, \partial_j U_i)^{1/2}$ the magnitude of the velocity gradient and Δ the IDDES length scale of OpenFOAM, $\Delta = \min\{\max[C_w d, C_w h_{\max}, h_{wn}], h_{\max}\}$ with $C_w = 0.15$, $h_{\max}$ the largest cell-centre-to-face distance and h_{wn} the wall-normal cell size. The step function equals one for $d \lesssim 0.5\Delta$ and vanishes beyond $d \approx 1.25\Delta$, so it guarantees the Inner state in the first near-wall cells whatever the value of r_d; it is the grid-dependent element of the shielding, kept in all campaigns as in IDDES. The second branch acts as a separation sensor based on the vorticity-Reynolds number,

$$l_\Omega = \sqrt{\frac{\nu_{\text{eff}}}{\max\left(\Omega, r_{vc}\nu_{\text{eff}}/d^2\right)}}, \tag{8}$$

$$x_\Omega = \frac{d}{l_\Omega} = \sqrt{\max\left(\text{Re}_\Omega, r_{vc}\right)}, \tag{9}$$

$$\text{Re}_\Omega = \frac{\Omega d^2}{\nu_{\text{eff}}}, \tag{10}$$

$$f_{d\Omega} = 1 - \tanh\left((c_{fd} x_\Omega)^3\right). \tag{11}$$

The Inner-Outer transition of this branch lies where $f_{d\Omega} = 1/2$,

$$\text{Re}_\Omega^* = \frac{[\text{artanh}(1/2)]^{2/3}}{c_{fd}^2} \approx \left(\frac{0.82}{c_{fd}}\right)^2, \tag{12}$$

i.e. $\text{Re}_\Omega^* \approx 67$ for $c_{fd} = 0.1$, an order of magnitude above the logarithmic-layer value $\text{Re}_\Omega = 1/\kappa^2 \approx 6$ (from $\Omega = u_\tau/\kappa d$ and $\nu_{\text{eff}} \approx \kappa u_\tau d$), so that attached boundary layers remain on the Inner set.

Throughout the present work, $c_{fd} = 0.1$, $c_d = 8.0$, and $r_{vc} = 1$ unless stated otherwise. The value $c_d = 8$ is the DDES constant of Spalart et al (2006). The constant c_{fd}, the single parameter that sets Re_Ω^*, is our choice, not a literature value, on two grounds. First, $\text{Re}_\Omega^* \approx 67$ keeps the whole attached boundary layer, logarithmic layer included, on the Inner set and selects the Outer set only in detached and free shear layers, where Re_Ω is much larger. Second, with $r_{vc} = 1$ the floor in l_Ω gives $f_{d\Omega} = 1 - \tanh(10^{-3}) \approx 1$ in quiescent, irrotational regions, so that this branch does not by itself switch such regions to the Outer set. Because the floor enters only through the product $c_{fd}\sqrt{r_{vc}}$, c_{fd} is held fixed in all campaigns and only the floor (r_{vc}, and its topology-sensitive extension of Section 2.2) is varied, which keeps the floor parameter identifiable.

2.2 Topology-sensitive extension

Preliminary calibrations revealed a persistent conflict between separated-flow and secondary-motion predictions. To test whether part of this deficit originates in the shielding itself, the constant floor r_{vc} is made topology-dependent,

$$r_{vc}(\mathbf{x}) = \max\left(r_{vc0}, C_{rvc}\frac{|\Omega|}{|\Omega| + |S|}\right), \tag{13}$$

where

$$|S| = \sqrt{2S_{ij}S_{ij}}, \qquad |\Omega| = \sqrt{2W_{ij}W_{ij}}. \tag{14}$$

The ratio

$$r = \frac{|\Omega|}{|\Omega| + |S|} \tag{15}$$

is a classical measure of local flow topology: $r \to 0$ in strain-dominated regions, $r = 1/2$ in a parallel shear flow and $r \to 1$ in rotation-dominated regions. The same balance of rotation and strain underlies the Ω-type vortex identification criteria (Liu et al, 2016) and the rotation measures of the SA rotation/curvature correction (Shur et al, 2000). The modification therefore raises the quiescent floor selectively in vortical regions and reverts to the constant floor elsewhere.

The lower bound $r_{vc0} = 1$ keeps the denominator of l_Ω from vanishing in irrotational regions. The gain C_{rvc} is not prescribed but calibrated as a seventh variable, searched in $[0.25, 400]$ alongside the six blending coefficients (Section 4.4); the two selected members settle at $C_{rvc} = 36$ (GSA-E) and $C_{rvc} = 0.31$ (GSA-N), i.e. the topology floor is active on the ERCOFTAC suite and inactive on the NASA suite ($C_{rvc}\, r \leq 0.31 < r_{vc0}$). With $C_{rvc} = 36$ the floor is ≈ 18 in a parallel shear flow ($r = 1/2$), above the logarithmic-layer value $\mathrm{Re}_\Omega \approx 6$, so that $f_{d\Omega} \approx 0.92$ there: the topology floor of GSA-E blends about 8% of the Outer set into the log layer. The topology floor is thus a controlled extension of the baseline shielding, not a replacement.

2.3 Nonlinear constitutive extension

Coefficient blending leaves the model a linear eddy-viscosity closure, bound by the Boussinesq approximation. To supply the anisotropy that drives secondary motion in corner flows, we add the optional QCR2024 nonlinear constitutive relation of Tamaki and Kawai (2024),

$$\tau_{ij}^{\mathrm{QCR}} = \tau_{ij} - C_{cr1}\left(O_{ik}\tau_{jk} + O_{jk}\tau_{ik}\right) - C_{cr2}\nu_t\frac{\mathrm{dev}\left(S_{ik}^{*}S_{kj}^{*}\right)}{|\nabla U|} - C_{cr3}\nu_t\sqrt{2W_{mn}W_{mn}}\delta_{ij}, \tag{16}$$

where

$$O_{ij} = \frac{2W_{ij}}{\max(|\nabla U|, \varepsilon)} \tag{17}$$

is the normalised rotation tensor, with $S_{ij} = (\partial U_i/\partial x_j + \partial U_j/\partial x_i)/2$, $W_{ij} = (\partial U_i/\partial x_j - \partial U_j/\partial x_i)/2$, S^*_{ij} the deviatoric part of S_{ij} and $\tau_{ij} = 2\nu_t S^*_{ij}$ the linear turbulent stress; the modelled Reynolds stress is $R_{ij} = \overline{u'_i u'_j} = -\tau^{\mathrm{QCR}}_{ij}$. The three terms are the original rotational QCR correction, a quadratic strain-based anisotropy term, and an isotropic term largely absorbed into the modified pressure in incompressible flow. The calibration objectives are therefore only weakly sensitive to C_{cr3}: with the two anisotropic terms switched off and $C_{cr3} = 2.15$, the velocity and shear-stress terms change by at most 4%, the wall-pressure terms of the curved step and of the gentle hill (the latter part of the search objective only) somewhat more, so C_{cr3} keeps its published value. The published SA-QCR2024 coefficients are [1]

$$(C_{cr1}, C_{cr2}, C_{cr3}) = (0.5, 1.0, 2.15). \tag{18}$$

Because the quadratic stress enters the momentum equation as an explicit source, its coupling with the segregated pressure-velocity algorithm can become unstable on external-aerodynamic meshes with strongly stretched cells while the solution is far from equilibrium. The implementation therefore provides two optional stabilisation parameters: a pseudo-time relaxation of the quadratic stress field, $\tau^{\mathrm{QCR,n}} \leftarrow \alpha_{\mathrm{QCR}}\,\tau^{\mathrm{QCR,n}} + (1-\alpha_{\mathrm{QCR}})\,\tau^{\mathrm{QCR,n-1}}$ with $\alpha_{\mathrm{QCR}} \in (0, 1]$, which alters only the convergence path, not the converged solution; and a realisability-motivated cap that limits the magnitude of the deviatoric (anisotropic) part of the quadratic stress to a multiple λ_{QCR} of the local linear stress magnitude, leaving the isotropic part as published. The cap modifies the converged state only where it clips, and it also removes the freestream indeterminacy of the $|\nabla U|$ normalisation, since the linear stress vanishes there. Both are inactive in every calibration run and are engaged only where QCR2024 is applied to external flows (Sections 4.5 and 4.9).

3 Calibration methodology

The Inner and Outer coefficient sets, the topology gain and the QCR constants add degrees of freedom, and a gain on one flow class may come with a loss on another. Coefficient selection is therefore posed as a multi-objective optimisation over suites of benchmark flows, not as case-by-case tuning. As sketched in Fig. 1, an automated framework couples steady RANS simulations in OpenFOAM-13 (Weller et al, 1998) with population-based optimisers, scores every candidate set on all cases of a suite relative to baseline SA and logs every error term of every evaluation.

3.1 Objective functions and optimisers

The objectives are built on three levels: error term, case and suite. For each case i, every error term m of Table 1 (a wall distribution such as C_p or C_f, or a set of profiles such as U or $\overline{u'v'}$) compares the sampled model quantity y, interpolated onto the

[1] The implementation permits Inner/Outer blending of the QCR coefficients, but they are kept globally uniform here to reduce parameter degeneracy and overfitting on few calibration cases (blending ablation in Section 4.5).

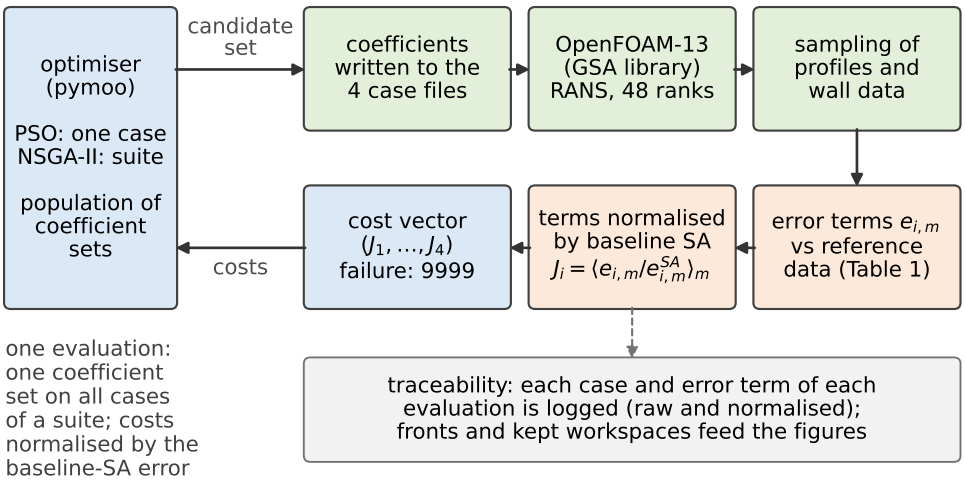

Fig. 1 The calibration loop. The optimiser proposes a population of coefficient sets; each set is written into the case templates of a suite, the four RANS cases are run with the GSA library, the sampled profiles and wall data are compared with the reference data (Table 1) and normalised by the baseline-SA error, and the resulting cost vector is returned to the optimiser.

reference abscissa, with the reference data y_{ref} through the relative mean-square error

$$e_{i,m} = \frac{1}{x_{\max} - x_{\min}} \int_{x_{\min}}^{x_{\max}} \left(\frac{y_{\mathrm{ref}}(x) - y(x)}{y_{\mathrm{n}}} \right)^2 \mathrm{d}x, \tag{19}$$

where y_{n} is the maximum of the reference quantity, or its maximum modulus for sign-changing quantities such as the shear stresses; terms defined at several stations sum $e_{i,m}$ over the stations. The hill shear-stress terms carry a guard against oscillatory profiles: at a station where the second derivative of the computed $\overline{u'v'}(y)$ changes sign more than five times with a magnitude above a fixed threshold, the station error is replaced by the number of sign changes, a penalty far above the error of any smooth profile. The integral is evaluated on the reference points with the composite Simpson rule for non-uniform spacing, exact for quadratic integrands; with an even number of points, the last interval uses the parabola through the last three points. The reported hill costs use velocity and shear stress only, without wall pressure, because the pressure level of a periodic domain is arbitrary; the gentle-hill wall-pressure term, part of the search objective, is removed from the stored terms.

Each term is divided by the same error of baseline SA, and the case cost is the mean of the N_i normalised terms of the case,

$$v_{i,m} = \frac{e_{i,m}}{\max\left(e^{SA}_{i,m}, \epsilon\right)}, \qquad J_i = \frac{1}{N_i} \sum_m v_{i,m}. \tag{20}$$

A case cost below unity is an improvement on baseline SA and above unity a degradation. The floor ϵ keeps physically insignificant variations from dominating where baseline SA is already near-exact. On the NASA suite $\epsilon = 10^{-2}$, a root-mean-square (rms) discrepancy of about 10% of the reference maximum, active on the terms marked with a dagger in Table 1 (flat plate, channel, hump velocity profiles): there the calibration is asked to preserve accuracy, not to improve it, and a cost of order 10^{-2} means "accuracy preserved", not "fifty times more accurate than baseline SA". On the ERCOFTAC suite $\epsilon = 0$. Baseline SA therefore has unit cost on every case except where the floor is active.

At suite level, the multi-case search treats the four case costs $(J_1, \ldots, J_4)$ as separate objectives, so that a degradation on one flow cannot be hidden by a gain on another, and returns a Pareto front of members none of which can be improved on one case without degrading another. Scalar measures, such as the mean cost $\bar{J} = \frac{1}{4}\sum_i J_i$ and the worst cost $\max_i J_i$, serve only to select members from a front (Section 3.3). The search is unconstrained, so front members, selected ones included, may exceed unity on a case.

Each case is first optimised alone by PSO, which minimises its cost J_i with 48 particles over 40 generations. The multi-case search uses the Non-dominated Sorting Genetic Algorithm II (NSGA-II) (Deb et al, 2002) with a population of 48, initialised with the single-case optima of the suite plus Latin-hypercube samples and evolved over warm-started rounds of up to 100 generations. Both optimisers are the `pymoo` implementations with a fixed random seed, and each population is evaluated in parallel on one compute node (24 to 48 cores); the retained runs cost about 9.9×10^3 core-hours in total. Every front reported is the final population of 48 non-dominated members, from a single realisation of the stochastic search, so differences of a few per cent between fronts are not interpreted. Diverged or failed evaluations receive a sentinel cost and are excluded from all statistics, and hill shear-stress profiles that oscillate in space are penalised, which raises the case cost several-fold.

3.2 Benchmark suites

The NASA suite comprises four cases of the TMR (Rumsey et al, 2010; Rumsey, 2026), referred to by their TMR identifiers: the zero-pressure-gradient flat plate (2DZP), the fully developed channel flow at high Reynolds number (2DFDC), the axisymmetric subsonic jet (ASJ) and the wall-mounted hump (2DWMH). They probe canonical near-wall behaviour, equilibrium shear flow, free-shear mixing and adverse-pressure-gradient separation.

The ERCOFTAC suite comprises the periodic hill (Breuer et al, 2009) in two configurations of the parametric DNS database of Xiao et al (2020) at $\mathrm{Re}_H = 5600$, in its 2021 extension that also varies the domain length and height: a steep hill ($\alpha = 0.5$, hill_{05} in the tables), whose large bubble is fixed by the geometry, and a gentle one ($\alpha = 1.5$, hill_{15}), with a thin pressure-induced bubble. The square duct at $\mathrm{Re}_\tau = 180$ (Vinuesa et al, 2014) and the curved backward-facing step (CBFS) (Bentaleb et al, 2012) complete the suite. These cases contain mechanisms that challenge eddy-viscosity models (separation, reattachment, curvature and turbulence-driven secondary motion) and form the core of the curated dataset of McConkey et al (2021)

and of the ERCOFTAC data-driven turbulence-modelling challenge (Buchanan et al, 2026; McConkey et al, 2026), which makes the present results directly comparable with data-driven closures trained on the same flows.

Table 1 lists the reference data and the error terms of every case. For 2DFDC the reference is the Spalding law with $\kappa = 0.40$ and $B = 5.5$, whereas SA is calibrated with $\kappa = 0.41$, so part of the baseline-SA channel error is the offset between the two logarithmic laws (Section 4.6).

Table 1 Reference data and error terms of the eight benchmark cases (†: term at the error floor, Section 3.1).

case	reference	error terms
2DZP	$C_f(\mathrm{Re}_\theta)$ correlation of the TMR (Rumsey, 2026)	C_f vs Re_θ (wall shear + momentum-thickness integral)†
2DFDC	Spalding law (κ=0.40, B=5.5)	$u^+(y^+)$ (u_τ from wall shear)†
ASJ	experiment (Bridges and Wernet, 2010)	$U(r)$, $\overline{u'v'}(r)$ at $x/D = 15, 20$
2DWMH	Greenblatt et al (2006) experiment	C_p, C_f, U and $\overline{u'v'}$ profiles at $x/c = 0.9, 1.0, 1.1$ ($U^\dagger$)
steep hill α=0.5	DNS (Xiao et al, 2020), $L_x \times L_y = 10.07h \times 3.04h$	U, $\overline{u'v'}$ at 7 stations x/h = 0.5-6 (+ oscillation penalty on $\overline{u'v'}$)
gentle hill α=1.5	DNS (Xiao et al, 2020), $L_x \times L_y = 10.93h \times 2.02h$	U, $\overline{u'v'}$ at 7 stations (+ oscillation penalty on $\overline{u'v'}$)
square duct	DNS (Re_τ=180) (Vinuesa et al, 2014)	U (4 lines), $\overline{u'v'}$ (2), $\overline{u'w'}$ (2)
CBFS	LES (Bentaleb et al, 2012)	C_p, C_f, U and $\overline{u'v'}$ profiles at $x/H = 2.5, 3.25, 4$

3.3 Calibration campaigns and model variants

The calibration is staged as in Fig. 2, independently on the two suites. A baseline-SA evaluation provides the normalising errors $e^{SA}_{i,m}$ of Eq. (20). The single-case PSO searches give a practical reference for per-case accuracy and the seeds of the multi-case search. Four multi-case NSGA-II campaigns follow, each adding one ingredient, with the search boxes of Table 2. The first, run1, calibrates the six blending coefficients $(\sigma_{\tilde{\nu}}, C_{b1}, C_{b2})_{I,O}$ in a box that brackets the baseline-SA values; run2 repeats the search in a widened box; run3 adds the topology gain C_{rvc} as a seventh variable, so that run2 is its constant-r_{vc} ablation reference; run4, on the ERCOFTAC suite only, freezes the blending at the run3 member GSA-E and searches C_{cr1} and C_{cr2}, tied Inner = Outer, over 26 generations seeded with the published constants, with $C_{cr3} = 2.15$ (Section 2.3). The staging separates the gains due to coefficient blending, shielding topology and nonlinear stress anisotropy.

The run1 box keeps $C_{b1} \geq 0$ and $C_{b2} \geq -1$, hence a non-negative front-spreading coefficient D (Section 2) and, through Eq. (3), a non-negative C_{w1}. The SA terms then keep the meaning they have in the original model (Spalart and Allmaras, 1992): the diffusion term propagates an eddy-viscosity front into non-turbulent fluid and the wall term destroys eddy viscosity. Within this box several optima lie on the $\sigma_{\tilde{\nu}}$

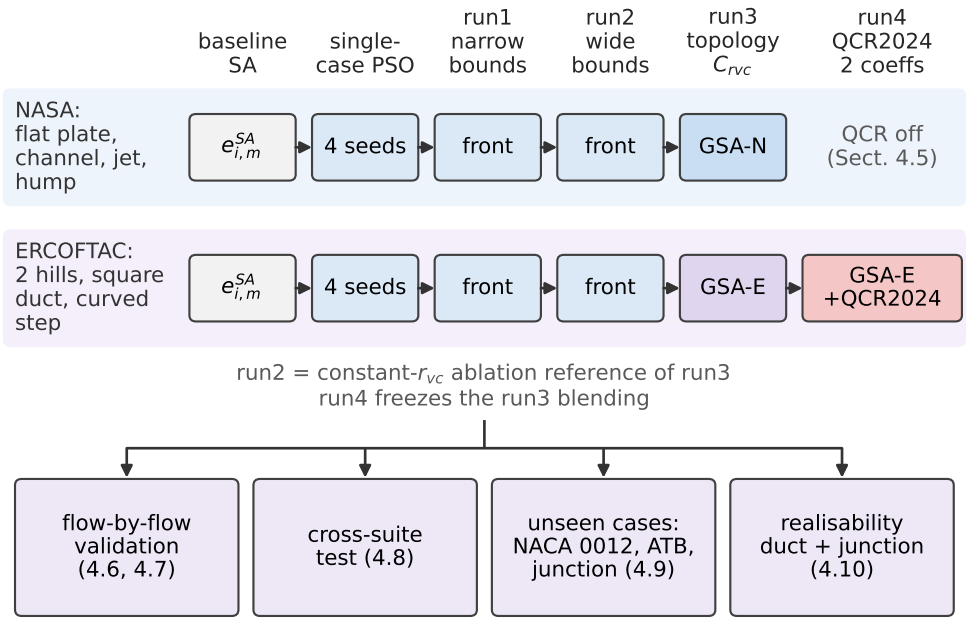

Fig. 2 Staged calibration strategy on the two suites and assessment chain. The run3 fronts deliver the selected members GSA-N (NASA) and GSA-E (ERCOFTAC); run4 adds the calibrated QCR2024 closure to GSA-E (GSA-E+QCR2024). We then compare the three members and baseline SA flow by flow, across suites, on three unseen configurations, and in the anisotropy-invariant plane.

Table 2 Search boxes of the multi-case campaigns, with the baseline-SA and published SA-QCR2024 (Tamaki and Kawai, 2024) values for reference. The single-case searches use the run3 box; run4 freezes the blending coefficients and C_{rvc} at GSA-E. The last row states whether the box guarantees the admissibility conditions $D \geq 0$ and $C_{w1} \geq 0$ of the original model.

coefficient	reference	run1	run2	run3	run4
$\sigma_{\tilde{\nu}\,I/O}$	2/3	$[0.5, 0.8]$	$[0.4, 1.0]$	$[0.4, 1.0]$	GSA-E
$C_{b1\,I/O}$	0.1355	$[0, 1]$	$[0, 1]$	$[0, 1]$	GSA-E
$C_{b2\,I/O}$	0.622	$[-1, 1]$	$[-2, 2]$	$[-2, 2]$	GSA-E
C_{rvc}	-	-	-	$[0.25, 400]$	GSA-E
C_{cr1}	0.5	-	-	-	$[0, 1]$
C_{cr2}	1.0	-	-	-	$[0, 2]$
C_{cr3}	2.15	-	-	-	2.15 (fixed)
$D \geq 0$, $C_{w1} \geq 0$ guaranteed	yes	yes	no	no	-

and C_{b2} bounds, and the ERCOFTAC members that improve every case do so only marginally (Section 4.2). The widened box of run2 and run3 is therefore a deliberate step beyond admissibility, taken to explore how far the SA structure can be pushed. It admits $D < 0$, for which the diffusion term no longer propagates an eddy-viscosity front, and $C_{w1} < 0$, for which wall destruction becomes a source. Both occur in the selected members ($D_I < 0$ for GSA-N and GSA-E, $C_{w1I} < 0$ for GSA-E, Section 4.3) and act wherever $\mathcal{F}_{SA} \approx 1$, which includes the irrotational freestream and the bulk of

internal flows (Section 2.1). This is a conscious choice. The calibrated sets are effective parameters within the envelope of the tested flows, not physically interpretable SA constants, and their gain over the admissible run1 measures how much accuracy the SA structure can deliver once its coefficients are freed from their original meaning.

The members carried through the paper are named after the suite that calibrated them. GSA-N and GSA-E are the run3 members of minimum mean cost $\bar{J}$ on the NASA and ERCOFTAC suites (Section 4.4). On NASA all four cases can be improved together; on ERCOFTAC, where the gentle hill and the duct conflict, this rule selects a member that improves on baseline SA in both. GSA-E+QCR2024 is the run4 member of minimum $\max(J_{\mathrm{hill}_{15}}, J_{\mathrm{duct}})$, which on this front is also the member of minimum worst cost $\max_i J_i$. The run1 and run2 compromises are the members of minimum mean cost on the NASA suite and, on the ERCOFTAC suite, the members of lowest duct cost among those that keep both hills and the step below unity. All variants of Table 3 are evaluated with the calibration pipeline from retained workspaces, so every difference between them is due to the coefficients alone. Line plots show only the variants of Table 3; the run1 and run2 compromises, omitted there for legibility, appear in the tables and in the coefficient and Pareto figures of Sections 4.1-4.4; field maps show baseline SA and the member calibrated on the suite, with the characteristic iso-line of every variant in one panel.

Table 3 Model variants compared in the flow-by-flow figures. Each variant keeps the same label, colour and line style in every figure of the paper.

label	definition
baseline SA	original coefficients (normalisation reference: unit cost, except for the terms normalised by the error floor, Table 4)
single-case optimum	the PSO optimum of that case alone (seven variables, as in run3)
GSA-N	NASA-calibrated member selected from the run3 front (Section 4.4)
GSA-E	ERCOFTAC-calibrated member selected from the run3 front (Section 4.4)
GSA-E+QCR2024	GSA-E with C_{cr1}, C_{cr2} calibrated in run4 and $C_{cr3} = 2.15$ fixed (Section 4.5)
GSA-E / GSA-N (cross-suite)	GSA-E on NASA, GSA-N on ERCOFTAC (Table 4, Section 4.8)
reference data	experiment, DNS, LES or analytical law of the case (Table 1)

4 Results

The results follow the calibration strategy. The single-case optima (Section 4.1) set the scale of what the generalised model can achieve on each flow and, through their coefficients, show which mechanisms each flow requires and where the flows conflict. The four multi-case campaigns follow (Sections 4.2-4.5), each read against the previous one. The calibrated models are then validated flow by flow on the NASA (Section 4.6) and ERCOFTAC (Section 4.7) suites, cross-validated on the other suite (Section 4.8) and tested on three unseen configurations (Section 4.9) with baseline SA and the three selected members only. The realisability analysis (Section 4.10) comes last because it uses both the ERCOFTAC duct and the unseen junction.

4.1 Single-case optima and their transfer to the other cases

Table 4 lists the cost of each single-case optimum on every case of its suite, i.e. what the generalised model achieves on the case it was tuned to and what this costs on the others. Table 5 lists the coefficients behind each gain, and Fig. 3 renders both tables graphically. Two observations follow from Table 4. First, the attainable single-case accuracy is high. The Inner/Outer split drives the flat-plate and channel errors to the floor, reduces the jet error by two orders of magnitude (0.013), halves the hump error (0.43) and reaches 0.47 and 0.15 on the gentle hill and the curved step. The exception is the square duct (0.80), where even a duct-dedicated calibration of a linear eddy-viscosity model removes only 20% of the error, which points to a structural deficit (Section 4.5). Second, single-case tuning is expensive: the flat-plate optimum multiplies the jet error by 21, the jet optimum raises the hump cost to 10.7, the step optimum laminarises the duct (zero modelled Reynolds stress) and raises its error fifty-fold, and the channel optimum makes the jet diverge. The single-case optima are therefore references, not usable models.

Table 5 explains why the per-case gains are large and why they cannot simply be added up. (i) The Inner and Outer states are used as two different models: in seven of the eight optima (all but the steep hill) the two values of at least one coefficient differ by more than half of its box. The channel and the step require a small Inner and a large Outer production, the jet and the flat plate the reverse. This is possible because the C_{w1} reconstruction of Eq. (3) keeps the log law intact for any Inner set. (ii) The flows disagree on the Outer state, which governs separation and mixing. The jet, a free shear layer, asks for less production than baseline SA ($C_{b1O} = 0.12$) at nearly baseline spreading ($D_O = 2.75$), because baseline SA over-mixes the jet (Section 4.6). The gentle hill ($\alpha = 1.5$) and the curved step, with pressure-induced separation from a curved wall, ask for three to seven times the baseline value ($C_{b1O} = 0.46$ and 0.98), because baseline SA under-mixes the separated shear layer and predicts bubbles that are too long. The hump sits in between ($C_{b1O} = 0.25$). The steep hill ($\alpha = 0.5$) is the exception. Its bubble is fixed by the steep lee slope and baseline SA already reproduces it closely, so more mixing would only shorten it, and its optimum halves the Outer production instead ($C_{b1O} = 0.07$). The two hills, one geometric family, therefore need opposite Outer states, and the conflict is asymmetric: the steep-hill optimum also improves the gentle hill (0.749), whereas the gentle-hill optimum degrades the steep one (1.495, Table 4). The shielding function selects the Outer state from local wall-distance and vorticity quantities, which do not distinguish a geometry-fixed from a pressure-induced separation, so a single Outer state cannot serve the jet and both hill shear layers at once. (iii) The sign of C_{b2} splits the cases. The flat plate, the jet and the gentle hill push C_{b2I} positive (a source at the edge of turbulent regions, faster spreading). The channel, the hump, the duct and the step drive C_{b2I} towards the lower bound (-1.6 to -1.8), where $D_I < 0$ turns the same term into a sink that thins the outer part of the attached boundary layer. Two optima need a caveat: the flat-plate optimum is degenerate (its cost is at the error floor, so any set that preserves the log law performs equally), and the duct optimum ($C_{b2} \approx -1.6$ in both states, $D < 0$) is the linear closure straining against a deficit it cannot represent. (iv) Every optimum except the channel one ($C_{rvc} = 0.8$) uses a large topology gain ($C_{rvc} = 32$-152), so the

topology floor of Section 2.2 is active whenever a single case is optimised; whether it survives in a multi-case compromise is tested by run3 (Section 4.4).

Table 4 Per-case normalised costs (Eq. (20)) of every variant on both suites: each single-case optimum evaluated on its full suite (in bold, the case it was tuned to), baseline SA, the multi-case compromises of run1-run4, and the selected member of the other suite (cross-suite rows, Section 4.8). QCR is off in every row except GSA-E+QCR2024. Baseline SA has unit cost except where the error floor of Section 3.1 applies (flat plate, channel and the velocity term of the hump). ‡: diverged.

NASA	flat plate	channel	jet	hump
baseline SA	0.022	0.035	1	0.899
optimum flat plate	**0.0002**	0.047	20.7	2.03
optimum channel	0.212	**0.005**	‡	4.52
optimum jet	1.354	0.046	**0.013**	10.7
optimum hump	1.039	0.023	5.68	**0.433**
run1 compromise	0.007	0.040	0.154	0.789
run2 compromise	0.017	0.037	0.108	0.676
GSA-N	0.063	0.040	0.048	0.715
GSA-E (cross-suite)	0.015	0.036	14.10	2.626
ERCOFTAC	hill_{05}	hill_{15}	duct	step
baseline SA	1	1	1	1
optimum hill_{05}	**0.776**	0.749	1.75	0.577
optimum hill_{15}	1.495	**0.467**	1.40	0.464
optimum duct	1.008	10.4	**0.802**	2.17
optimum step	1.386	0.496	51.3	**0.151**
run1 compromise	0.950	0.815	0.896	0.967
run2 compromise	0.864	0.710	1.198	0.779
GSA-E	1.072	0.625	0.920	0.308
GSA-E+QCR2024	0.947	0.584	0.138	0.422
GSA-N (cross-suite)	0.840	0.941	1.348	0.933

Table 5 Coefficients of the single-case optima, of the run1 and run2 compromises and of the selected run3 members GSA-N and GSA-E (Inner/Outer values; costs of every member in Table 4; C_{b2} and C_{rvc} of the selected members to the precision of the archived sets). $D = (1 + C_{b2})/\sigma_{\tilde{\nu}}$ is the front-spreading coefficient of the diffusion term and $C_{w1} = C_{b1}/\kappa^2 + D$ the reconstructed destruction coefficient (Eq. (3)), $\kappa = 0.41$); baseline SA has $D = 2.43$ and $C_{w1} = 3.24$. C_{rvc} is the gain of the topology floor (Eq. (13)).

member	$\sigma_{\tilde{\nu}\,I/O}$	$C_{b1\,I/O}$	$C_{b2\,I/O}$	$D_{I/O}$	$C_{w1\,I/O}$	C_{rvc}	own cost
baseline SA	0.667/0.667	0.136/0.136	+0.62/+0.62	2.43/2.43	3.24/3.24	-	1
NASA suite							
optimum flat plate	0.638/0.831	0.673/0.022	+1.12/−1.42	3.32/−0.51	7.33/−0.37	152	0.0002
optimum channel	0.992/0.402	0.064/0.660	−1.67/−0.05	−0.67/2.36	−0.29/6.29	0.8	0.005
optimum jet	0.420/0.417	0.852/0.118	+1.34/+0.15	5.58/2.75	10.65/3.45	104	0.013
optimum hump	0.582/0.992	0.047/0.245	−1.82/−0.35	−1.40/0.65	−1.12/2.11	83	0.433
run1 compromise	0.794/0.512	0.164/0.121	−0.99/+0.31	0.02/2.56	1.00/3.28	-	-
run2 compromise	0.718/0.408	0.170/0.106	−1.51/+1.80	−0.71/6.86	0.31/7.49	-	-
GSA-N	0.961/0.400	0.164/0.106	−1.445/+1.370	−0.46/5.92	0.51/6.55	0.31	-
ERCOFTAC suite							
optimum hill_{05}	0.999/0.999	0.257/0.069	−0.82/+0.82	0.18/1.82	1.71/2.23	83	0.776
optimum hill_{15}	0.401/0.980	0.092/0.459	+2.00/+2.00	7.49/3.06	8.03/5.79	125	0.467
optimum duct	0.406/0.771	0.123/0.163	−1.61/−1.56	−1.51/−0.73	−0.78/0.24	143	0.802
optimum step	0.995/0.403	0.021/0.982	−1.58/−2.00	−0.58/−2.48	−0.45/3.36	32	0.151
run1 compromise	0.735/0.514	0.109/0.244	−0.94/−0.97	0.08/0.06	0.73/1.51	-	-
run2 compromise	0.843/0.464	0.155/0.239	−1.39/−1.22	−0.46/−0.47	0.46/0.95	-	-
GSA-E	0.695/0.882	0.092/0.442	−1.563/+1.132	−0.81/2.42	−0.26/5.05	35.98	-

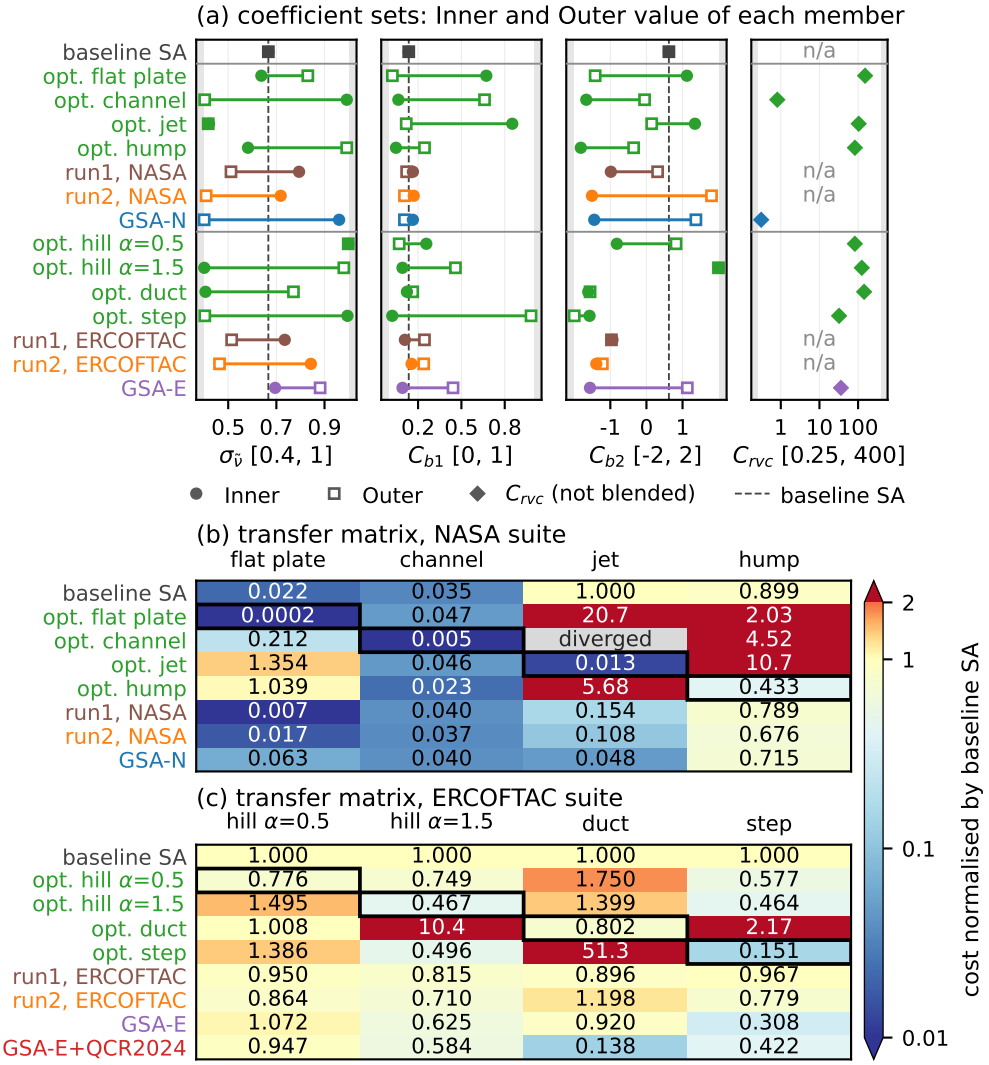

Fig. 3 Visual summary of Tables 4 and 5. (a) Inner (filled circle) and Outer (open square) values of $\sigma_{\tilde{\nu}}$, C_{b1} and C_{b2} of every member, joined by a bar whose length measures how differently the two blended states are used; each axis spans the optimiser box printed below it (grey margins: outside the box; dashed line: baseline SA). Last column: the topology gain C_{rvc}, one value per member (none for baseline SA and for the six-variable run1 and run2 members). (b), (c) Transfer matrices: normalised cost of every member on each case of its suite (cross-suite rows of Table 4 not drawn), on a logarithmic colour scale from 0.01 to 2 centred on baseline SA (unity), blue better, red worse, costs beyond the range in the extreme colours; black box: the case each single-case optimum was tuned to.

4.2 Narrow-bounds blending calibration (run1)

Campaign run1 searched the six blending coefficients within the narrow bounds of Section 3.3, inside which every set is physically admissible ($D \geq 0$ and $C_{w1} \geq 0$ in both states). On NASA, 3 of the 48 front members are already below unity on all four cases. The run1 compromise (mean cost 0.25) keeps the flat plate and the channel at the error floor and brings the jet to 0.154 and the hump to 0.789 (Table 4). The per-case front minima are lower (Table 6), so the front contains near-optimal members for every case, but the jet and the hump pull the Outer state in opposite directions (Section 4.1), and the compromise concedes accuracy on the hump. On ERCOFTAC, 3 members are below unity on all four cases (costs 0.77-0.98) and 6 on both the gentle hill and the duct. The run1 compromise improves every case, but only marginally (costs 0.81-0.97).

The run1 fronts press against the box. On NASA, 37 of the 48 members have $\sigma_{\tilde{\nu}I} > 0.79$, 35 have $\sigma_{\tilde{\nu}O} < 0.51$ and 29 have $C_{b2I} < -0.95$; on ERCOFTAC, 30 have $\sigma_{\tilde{\nu}O} < 0.51$. The optimiser thus asks for more Outer spreading and destruction than the box allows, and drives the Inner diffusion coefficient to the admissibility limit $D_I = 0$ itself ($D_I = 0.02$ and 0.08 for the two run1 compromises, Table 5). The widened bounds of run2 deliberately step past this limit.

Table 6 Front statistics of the three blending campaigns, over the finite members of each 48-member front; costs normalised by baseline SA. Campaigns run1: narrow bounds; run2: widened bounds, constant r_{vc}; run3: widened bounds, topology-sensitive r_{vc}. "Corner <1": members below unity on both hill_{15} and duct; "all four <1": below unity on every case of the suite; "best-hill duct": duct cost of the member with the lowest hill_{15} cost (that cost in brackets); "best mean": lowest mean cost over the four cases among the corner <1 members; "min-max": smallest $\max(\text{hill}_{15}, \text{duct})$.

NASA	per-case front minimum				all four <1
	flat plate	channel	jet	hump	
run1 (narrow bounds)	0.0024	0.0182	0.1321	0.4408	3/48
run2 (constant r_{vc})	0.0035	0.0065	0.0378	0.4452	6/48
run3 (topology r_{vc})	0.0038	0.0085	0.0479	0.4587	10/48
ERCOFTAC	corner <1	best-hill duct	best mean	min-max	all four <1
run1 (narrow bounds)	6/48	1.32 (0.507)	0.855	0.887	3/48
run2 (constant r_{vc})	3/48	51.3 (0.459)	0.765	0.941	0/48
run3 (topology r_{vc})	4/48	1.27 (0.470)	0.731	0.896	1/48

4.3 Widened-bounds blending calibration (run2)

Campaign run2 repeats the search with the widened bounds of Section 3.3. On NASA the wider box improves the front. The members below unity on all four cases grow from 3 to 6 of 48, and the run2 compromise lowers the mean cost from 0.25 to 0.21, the jet cost from 0.154 to 0.108 and the cost of the hump, the historically hard separation case, from 0.789 to 0.676 (Table 4). The flat plate and the channel stay at the error floor. The compromise uses the Inner/Outer split as designed (Table 5): near-baseline Inner production, the low Outer production of the jet optimum and a strong Outer spreading term ($D_O = 6.9$).

On ERCOFTAC the wider box extends the reach of the front on single cases (front minima: gentle hill 0.507 → 0.459, duct 0.80 → 0.78, step 0.192 → 0.161) and allows larger overall gains: among the members below unity on both the gentle hill and the duct, the lowest mean cost falls from 0.855 to 0.765. The corner-flow trade-off, however, becomes sharper. The gentle hill ($\alpha = 1.5$, over-separated by baseline SA) and the square duct pull in opposite directions, so that no member is below unity on all four cases and only 3 of 48 are below unity on both corner cases (6 in run1). The run2 compromise keeps both hills and the step below unity, with the duct at 1.20. The corner secondary flow requires Reynolds-stress anisotropy that no linear eddy-viscosity model reproduces, with or without blending (Demuren and Rodi, 1984), a deficit addressed by the topology floor (run3) and QCR2024 (run4). Figure 4b places the run2 and run3 fronts in the gentle-hill/square-duct plane, where every run2 member that improves on baseline SA on both corner cases is marginal on at least one of them (cost above 0.94).

The widened box also takes the compromises out of the admissible range of run1. Every widened-bounds compromise has $C_{b2I} < -1$, hence a negative Inner diffusion coefficient ($D_I = -0.81$ to -0.46; the ERCOFTAC run2 compromise also has $D_O < 0$), and GSA-E, like five of the eight single-case optima, also has a negative destruction coefficient in one state ($C_{w1I} = -0.26$, Table 5): the diffusion and wall terms lose their original meaning, and the coefficients act as effective parameters that reshape the eddy viscosity of the tested flows (Section 3.3).

4.4 Topology-sensitive r_{vc} calibration (run3)

The topology floor of Section 2.2 raises x_Ω, and hence lowers $f_{d\Omega}$, wherever it exceeds the local Re_Ω: in rotation-dominated regions of low Re_Ω, such as vortex cores and separated shear layers, the Outer coefficients then act where the constant floor kept the Inner ones. Campaign run3 is a controlled ablation against run2 (same population, generations and bounds, plus C_{rvc}).

On the NASA suite the topology floor is nearly neutral. The per-case front minima are 3-31% above those of run2 (Table 6), differences too small to interpret from a single realisation per campaign (Section 3.1). The members below unity on all four cases grow from 6 to 10 of 48 (from 5 to 9 if the hump must also improve on the cost of baseline SA itself, 0.899). The selected GSA-N, the run3 member of minimum mean cost (0.217, against 0.210 for the run2 compromise), trades a higher flat-plate cost, still at the error floor, for a lower jet cost, with nearly unchanged channel and hump costs (Table 4). Its $C_{rvc} = 0.31$ leaves the topology floor inactive, so GSA-N is in practice a constant-r_{vc} model. Like the run2 compromise, it keeps the jet-like Outer production ($C_{b1O} = 0.11$) and compensates the hump with a strong spreading term ($D_O = 5.9$, $C_{w1O} = 6.6$), which acts in the near-wall separated layer, where the destruction term is active, but is inert in the jet. As shown in Fig. 4a, the run2 and run3 fronts occupy the same region of the jet-hump plane, and most members below unity on both the jet and the hump also keep the flat plate and the channel below unity (all 6 in run2, 10 of 13 in run3).

On the ERCOFTAC suite the topology floor relieves part of the corner-flow trade-off that the widened bounds sharpened. Four front statistics quantify the change with respect to run2 (Table 6). The members below unity on both corner cases grow from

3 to 4 of 48 (one of them, close to the step optimum, sits at the edge of duct laminarisation), and run3 recovers one member below unity on all four cases (none in run2, three marginal ones in run1). The duct cost of the best gentle-hill member falls from 51.3 (run2, gentle hill 0.459) to 1.27 (run3, gentle hill 0.470); 51.3 is the coefficient-independent cost of a laminarised duct solution with zero modelled Reynolds stress, also scored by the step optimum (Table 4). The balanced $\max(\text{hill}_{15}, \text{duct})$ improves from 0.941 to 0.896, close to the 0.887 of run1, while the lowest mean cost among the members below unity on both corner cases falls further, from 0.765 to 0.731. As shown in Fig. 4b, every run2 member of this group is dominated in the gentle-hill/square-duct plane by a run3 member.

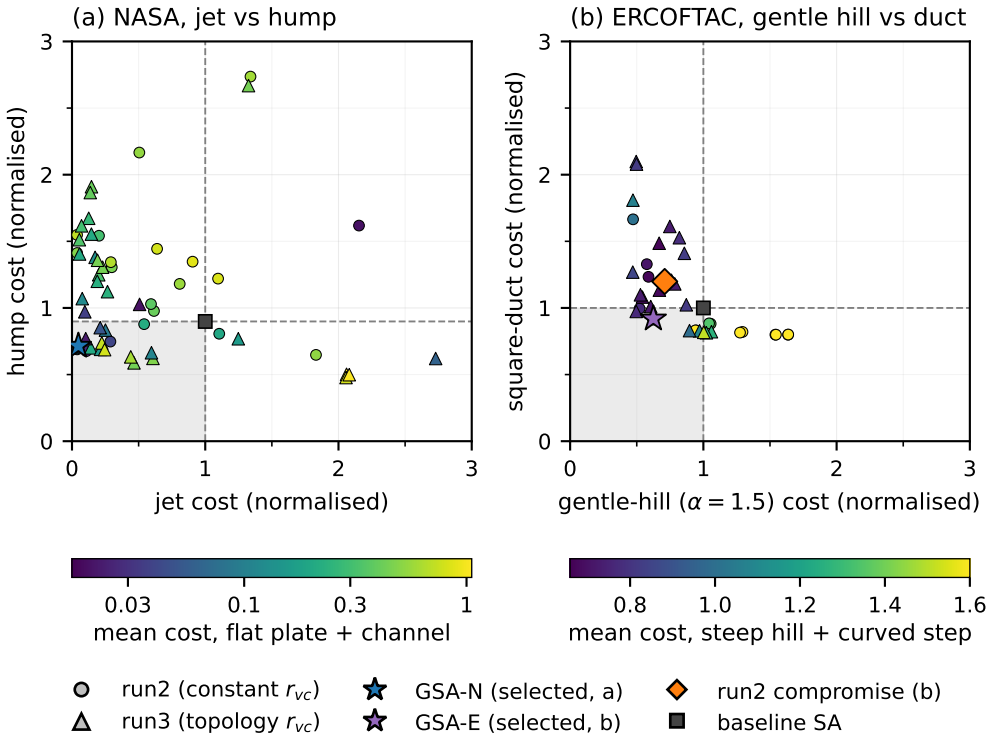

Fig. 4 Pareto fronts of run2 (circles, constant r_{vc}) and run3 (triangles, topology-sensitive r_{vc}) in the trade-off plane of each suite, coloured by the mean normalised cost of the two remaining cases of the suite (colour bar under each panel). Grey rectangle: better than baseline SA (square, at its own cost, 0.90 on the hump because of the error floor) on both plotted cases. (a) NASA, jet vs hump (colour: flat plate and channel), with GSA-N (star). (b) ERCOFTAC, gentle hill ($\alpha = 1.5$) vs square duct (colour: steep hill and curved step), with GSA-E (star) and the run2 compromise (diamond). Not shown: members with a diverged case (a: 11 run2, 3 run3; b: 4 run2, 1 run3) or a cost above 3 (a: 16 run2, 11 run3; b: 29 run2, 24 run3).

GSA-E, the run3 member of minimum mean cost (0.731), is carried forward, with its SA coefficients frozen, to the run4 QCR calibration. The only member below unity

on all four cases, which also attains the balanced $\max(\text{hill}_{15}, \text{duct})$ of 0.896, is not selected because its gains are small on every case (costs 0.830-0.998, mean 0.927), whereas GSA-E accepts a 7% loss on the steep hill (1.07, its only cost above unity) for much lower gentle-hill and step costs (0.62 and 0.31). GSA-E adopts the hill and step reading of the Outer state ($C_{b1O} = 0.44$, against 0.11 for GSA-N) with near-baseline spreading ($D_O = 2.42$, Table 5). This Outer production is most likely the main cause of the cross-suite asymmetry (Section 4.8) and of the stall behaviour (Section 4.9), although the two members also differ in C_{rvc} and in the Inner set. The topology floor thus relaxes the run2 trade-off, at no significant cost on the NASA suite, but does not close it, because no member places both corner cases comfortably below unity and the two hills still need opposite Outer states (Section 4.1). The run3 campaign therefore marks how much of the ERCOFTAC deficit the shielding function alone can reach, and motivates the QCR2024 stage.

4.5 Calibrated QCR2024 anisotropy (run4)

In run4 the blending coefficients are frozen at GSA-E and C_{cr1}, C_{cr2} are calibrated (Section 3.3); Fig. 5 compares the run4 front with the run3 front (a) and the per-case costs of the selected member with those of GSA-E (b).

QCR2024 acts mainly on the duct, through the balance of its rotation and strain terms. In a parallel shear flow the closure sets the wall-parallel minus the wall-normal cross-plane normal stress to $\nu_t|\nabla U|\,(2C_{cr1} - C_{cr2}/4)$ (Section 4.10), and the run4 front follows this combination: every member with $2C_{cr1} < C_{cr2}/4$ has a duct cost no lower than that of GSA-E (0.92-2.3), every member with $2C_{cr1} > C_{cr2}/4$ a lower one, down to 0.12. The rotation term thus drives the corner vortices, and the strain term alone degrades the duct. Across the 20 front members that pass the oscillation guard, the gentle hill stays within 7% of GSA-E (0.58-0.62), the steep hill falls from 1.07 to 0.95-1.04 and the step ranges from 0.29 to 0.42. The other 28 members, the four of lowest duct cost (0.124-0.133) among them, trip the guard on the gentle hill, as do the published constants (Table 7), all at the first station, $x/h = 0.5$, with six or seven sign changes against the five allowed, which raises the gentle-hill cost to 6.2-7.2. Without the guard their gentle-hill costs would be 0.58-0.62, as for the other members, and the three closest to GSA-E+QCR2024 ($C_{cr1} = 0.47$-0.48, $C_{cr2} = 1.52$-1.56) would match its balanced cost (0.583 against 0.584). The corner trade-off of Section 4.3 is nevertheless resolved: 10 run4 members improve on baseline SA on both corner cases and 9 on all four cases (at most 6 and 3 in run1-run3), and the balanced $\max(\text{hill}_{15}, \text{duct})$ falls to 0.584 (at best 0.887 before, Table 6). The min-max member, GSA-E+QCR2024, settles at $C_{cr1} = 0.45$ and $C_{cr2} = 1.50$, against the published 0.5 and 1.0; it is also the member of lowest duct cost among those that pass the guard. As listed in Table 7, GSA-E+QCR2024 lowers the duct cost by 85% ($0.920 \rightarrow 0.138$), lowers both hill costs (steep hill $1.072 \rightarrow 0.947$, now below unity; gentle hill $0.625 \rightarrow 0.584$) and raises the step cost from 0.308 to 0.422. Restricting the calibrated closure to one blending state shows where it acts: active only in the Inner state it retains the duct gain (duct cost 0.147, against 0.138 for the full closure) but trips the oscillation guard on both hills (costs 14.1 and 6.2); active only in the Outer state it keeps the hills and the step near the costs of GSA-E or below (0.955, 0.626 and

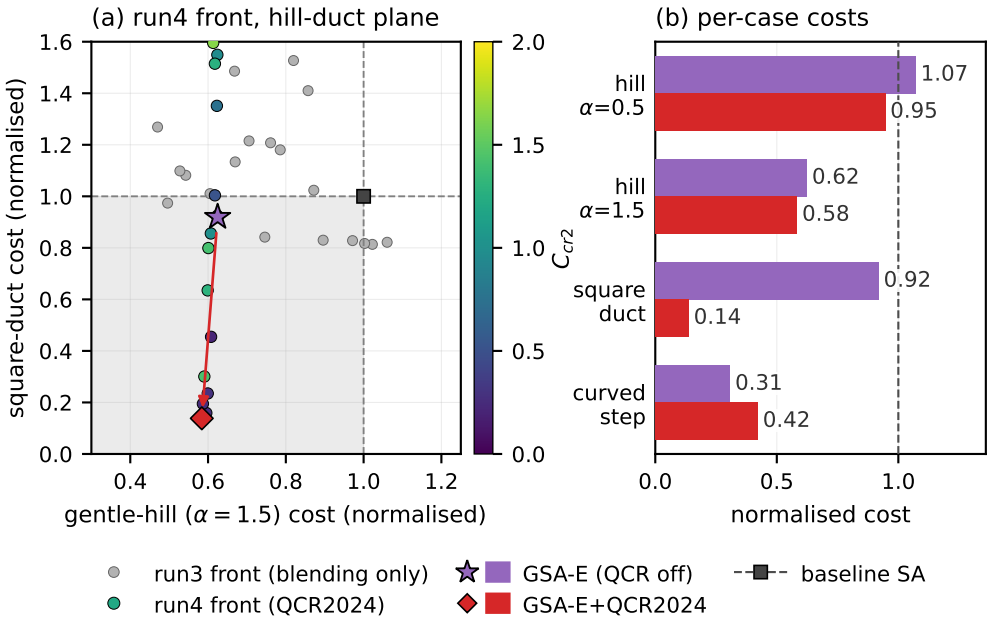

Fig. 5 QCR2024 calibration of run4 on the ERCOFTAC suite (blending frozen at GSA-E). (a) Gentle hill vs square duct: run3 front (grey; 28 members off scale), GSA-E (star), run4 front (coloured by C_{cr2}) and GSA-E+QCR2024 (diamond, arrow from GSA-E); grey rectangle: better than baseline SA (square) on both cases. The run4 members that pass the oscillation guard (Section 3.1) form a vertical band at or just below the gentle-hill cost of GSA-E; off scale: 5 of them with a duct cost above 1.6 and the 28 members that trip the guard on the gentle hill (cost above 6). (b) Per-case costs of GSA-E+QCR2024 and of GSA-E with QCR off; dashed line: baseline SA.

0.327) but recovers only a small part of the duct gain (0.810), because $\mathcal{F}_{SA} > 0.5$ in more than 99.5% of the duct cells (Section 2.1). Inner/Outer blending of the quadratic closure thus offers no additional leverage on the duct, and the QCR constants are kept global. The corresponding duct flow is shown in Section 4.7.

Table 7 QCR2024 calibration of run4 on the ERCOFTAC suite (blending frozen at GSA-E; costs normalised by baseline SA). *Published value, fixed (Section 2.3). †Oscillation guard of Section 3.1 tripped at $x/h = 0.5$; 0.582 without the guard.

member	C_{cr1}	C_{cr2}	C_{cr3}	hill_{05}	hill_{15}	duct	step
GSA-E, QCR off	0	0	0	1.072	0.625	0.920	0.308
published constants	0.5	1.0	2.15	0.951	6.221†	0.205	0.447
GSA-E+QCR2024 (min-max compromise)	0.451	1.496	2.15*	0.947	0.584	**0.138**	0.422

QCR2024 is not part of the NASA calibration, whose suite contains no corner or secondary flows, the flows for which it was designed, and switching it on in GSA-N is detrimental. Evaluated with the quadratic terms active (published constants in both or in only one blending state, or the run4 set in both, with the stabilised formulation of Section 2.3), GSA-N degrades on the two objectives governed by the Outer state:

the jet cost grows by about half (0.048 → 0.072-0.073) and the hump by 4-21% (0.715 → 0.75-0.87), while the flat plate and the channel stay at the error floor (0.055-0.062 and 0.040); with the closure active only in the Inner state the jet and the hump diverge. These four settings suggest that a NASA calibration of the closure would drive its constants towards zero; a joint recalibration of the blending and QCR2024 coefficients on external flows is left to future work.

4.6 NASA suite: flow-by-flow validation

On 2DZP all variants reproduce the $C_f(\mathrm{Re}_\theta)$ correlation to within 4%. As shown in the zoomed panel of Fig. 6b, the single-case optimum sits on top of the correlation (within 0.3%), baseline SA runs 0.5-3% low (1.4% on average over $4000 \le \mathrm{Re}_\theta \le 13000$), and GSA-N runs 1.6-3.7% high (2.9% on average). This excess is a real regression with respect to baseline SA, the price paid for the jet and hump gains, which the error floor accepts by design: the floor keeps the calibration from spending coefficient authority on C_f shifts of a few per cent, and the flat-plate optimum of Table 5 should not be read as physics.

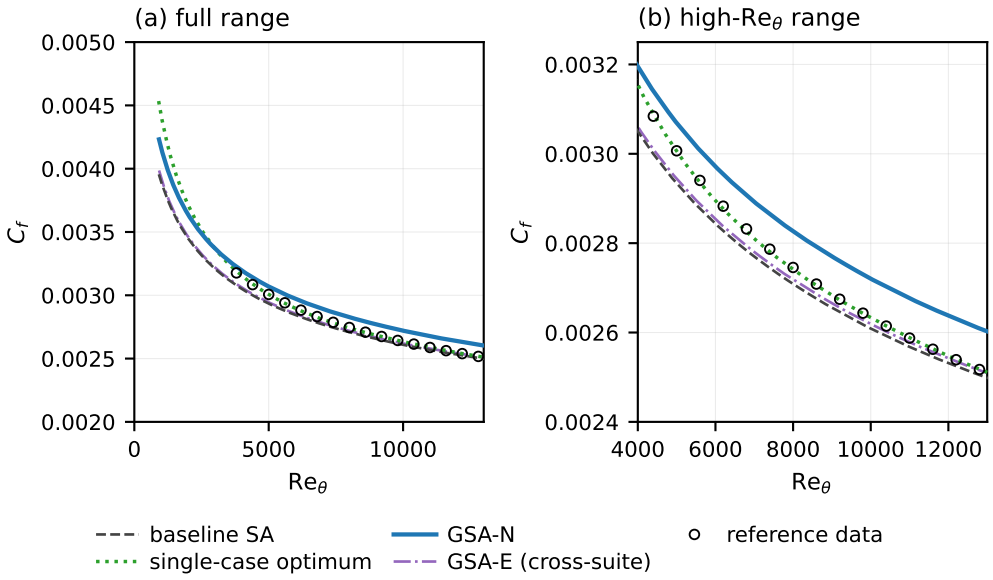

Fig. 6 2DZP: skin friction vs momentum-thickness Reynolds number for all variants (C_f from the wall shear, Re_θ from the momentum-thickness integral of the sampled profiles, as in the cost metric). (a) full range; (b) zoom on the range covered by the correlation ($4000 \le \mathrm{Re}_\theta \le 13000$).

In 2DFDC all variants follow the law of the wall to within $1.5\,u^+$ up to $y^+ = 10^5$. As shown in Fig. 7, the buffer layer overshoots the Spalding profile by up to $1.2\,u^+$ for every variant, and the C_{w1} reconstruction of Eq. (3) keeps the slope of the log layer, since between $y^+ = 10^3$ and 10^5 the deviation changes by less than $0.25\,u^+$.

Baseline SA follows its own logarithmic law (a fit over $10^3 < y^+ < 10^5$ gives $\kappa = 0.409$, $B = 4.9$) and therefore sits $\approx 1\,u^+$ ($\approx 3\%$ of the centreline velocity) below the Spalding reference from $y^+ \approx 10^3$ outward, an offset between the two sets of log-law constants (Section 3.2), not a model deficit. The channel optimum tracks the reference law across four decades of y^+ (within $0.7\,u^+$ up to $y^+ = 10^6$; its own log-layer fit gives $\kappa \approx 0.39$-0.40, i.e. it adopts the reference constants) by thinning the outer part of the profile through the negative Inner gradient-diffusion coefficient ($D_I < 0$, Table 5). The multi-case members, which retain a similar C_{b2I} but a higher Inner production, inherit the offset of baseline SA, and GSA-N enlarges it to $\approx 1.6\,u^+$ for $10^5 < y^+ < 10^6$ and $\approx 2.5\,u^+$ ($\approx 6\%$ of the centreline velocity) beyond $y^+ = 10^6$, a genuine regression in the outer part of the layer, again accepted by the error floor.

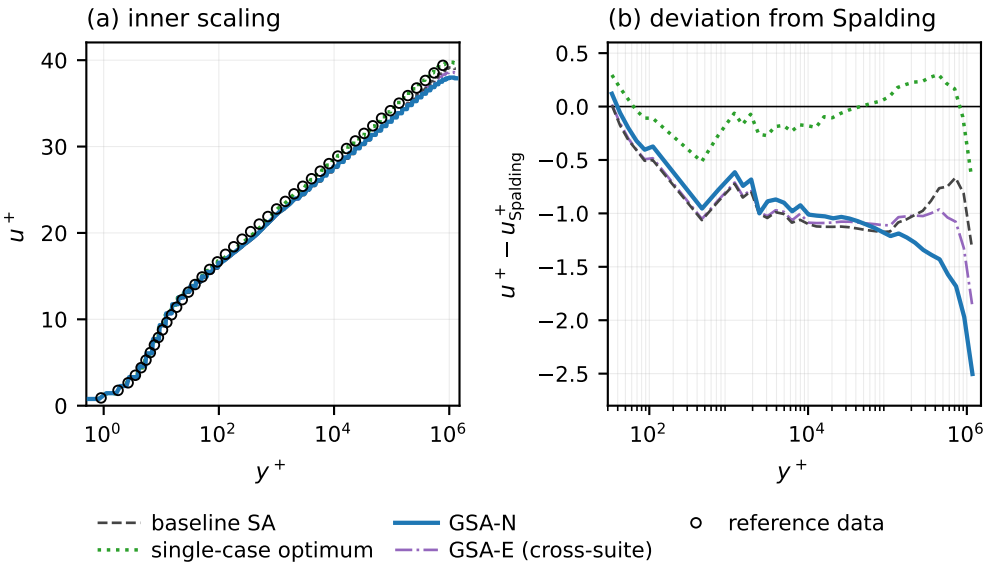

Fig. 7 2DFDC: (a) $u^+(y^+)$ against the Spalding law for all variants (u_τ from the wall shear at $x = 500$, as in the cost metric); (b) deviation from the law for $y^+ > 30$ (log-binned).

ASJ shows the largest single-case gain of the calibration and gives the clearest picture of the role of the Outer coefficient set. As shown in Fig. 8a, baseline SA ends the potential core slightly early ($U_c/U_{\mathrm{jet}} = 0.95$ at $x/D = 6.7$, against $x/D = 7.1$ in the TMR measurements (Rumsey, 2026)) and over-spreads the shear layer. Its centreline velocity therefore falls 17%, 24% and 25% below the measurement at $x/D = 10$, 15 and 20, as shown in Fig. 9a. The calibrated member GSA-N restores the measured core length ($x/D = 7.1$). In Fig. 9b,c it lies on top of the measured $U(r)$ profiles at all five stations $x/D = 2$-20 (rms error $\leq 2.5\%$ of the measured peak) and of the $\overline{u'v'}(r)$ profiles at the two stations that enter the cost, $x/D = 15$ and 20 (jet cost 0.048). The jet optimum is closer still to the measurements from $x/D = 10$ (rms error $\leq 2.7\%$ for U and $\leq 3.7\%$ for $\overline{u'v'}$), but its core is shorter ($x/D = 6.5$, with a centreline overshoot

to $U_c/U_{\text{jet}} = 1.03$) and it departs from the measurements upstream, at $x/D = 2$ and 5. The responsible coefficient is the Outer production, reduced from 0.1355 to 0.11-0.12 in the jet optimum, the run1 and run2 compromises and GSA-N (Table 5), because far from walls the production-diffusion balance alone sets the eddy viscosity (Section 2). A single Outer production tuned on a round jet need not suit other free shear flows: SA-type closures do not reproduce the different spreading rates of round and plane jets (the round-jet/plane-jet anomaly (Pope, 1978)), and the suites contain no plane jet, mixing layer or wake.

On 2DWMH the shielding keeps $\mathcal{F}_{SA} \approx 1$ (Inner state) in the attached boundary layer up to the crest and switches to the Outer state in the separated shear layer, as shown in Fig. 10d. The hump optimum shortens the recirculation bubble from that of baseline SA ($U_x = 0$ closure at $x/c \approx 1.14$, skin-friction reattachment at 1.17) to the reattachment measured by Greenblatt et al (2006) ($x/c \approx 1.11$). The gain is clearer in the skin friction than in the velocity field: the $U_x = 0$ closures of baseline SA and of the optimum differ by only $0.03c$ (Fig. 10e), whereas the C_f of the optimum (Fig. 11b) recovers downstream of reattachment at the measured rate (mean $C_f = 1.32 \times 10^{-3}$ over $1.2 < x/c < 1.3$, against 1.26×10^{-3} measured and 0.60×10^{-3} for baseline SA). Inside the bubble the optimum overpredicts the magnitude of the backflow skin friction (-1.15×10^{-3} against -0.54×10^{-3} measured over $0.7 < x/c < 0.9$). GSA-N keeps a bubble of essentially baseline-SA length ($U_x = 0$ closure at $x/c \approx 1.16$) and improves the case through the pressure and velocity fields instead. As shown in Fig. 11a, its C_p plateau over the bubble deepens towards the experiment, as does that of the hump optimum, and its post-reattachment pressure recovery moves onto the data. In Fig. 11c,d its separated-region velocity profiles follow the measured backflow and mixing-layer shape, and the $\overline{u'v'}$ peaks, underpredicted by about half by baseline SA in the bubble, rise towards the data.

Because this separated shear layer is bounded by the wall, the Outer destruction term is active, and the two calibrated sets reach their gains through different Outer coefficients (Table 5). The hump optimum (0.433) combines more Outer production than baseline SA ($C_{b1O} = 0.25$ against 0.1355) with less spreading ($D_O = 0.65$ against 2.43) and less destruction ($C_{w1O} = 2.11$ against 3.24). It therefore raises the eddy viscosity where the separated shear layer forms, as shown in Fig. 10a-c: the peak of ν_t/ν at $x/c \approx 0.9$ grows from 810 to 1100. The stronger mixing raises the $\overline{u'v'}$ peak at $x/c = 1.0$ from 0.014 to $0.018\,U_\infty^2$ (measured: $0.030\,U_\infty^2$), deepens the C_p plateau (mean offset from the data over $0.8 < x/c < 1.0$ reduced from 0.099 to 0.035) and brings reattachment, and with it the C_f recovery, onto the measurements. GSA-N (0.715) instead keeps the jet-like Outer production and relies on a strong spreading term (Section 4.4). Its eddy viscosity starts over a wider layer (width at half maximum $0.10c$ against $0.065c$ for baseline SA at $x/c \approx 0.8$) but rises later (900 at $x/c \approx 0.9$), and stays above that of baseline SA downstream of reattachment. This improves the velocity profiles at the calibration stations (rms deviation at $x/c = 1.0$ reduced from 0.045 to $0.023\,U_\infty$) and halves the C_p error of the post-reattachment recovery ($1.0 < x/c < 1.3$), but not the skin friction. The C_f of GSA-N crosses zero only at $x/c \approx 1.21$ and reaches 0.22×10^{-3} over $1.2 < x/c < 1.3$, a slower recovery than that of baseline SA. On this flow, therefore, C_f follows how early the Outer production raises the eddy viscosity

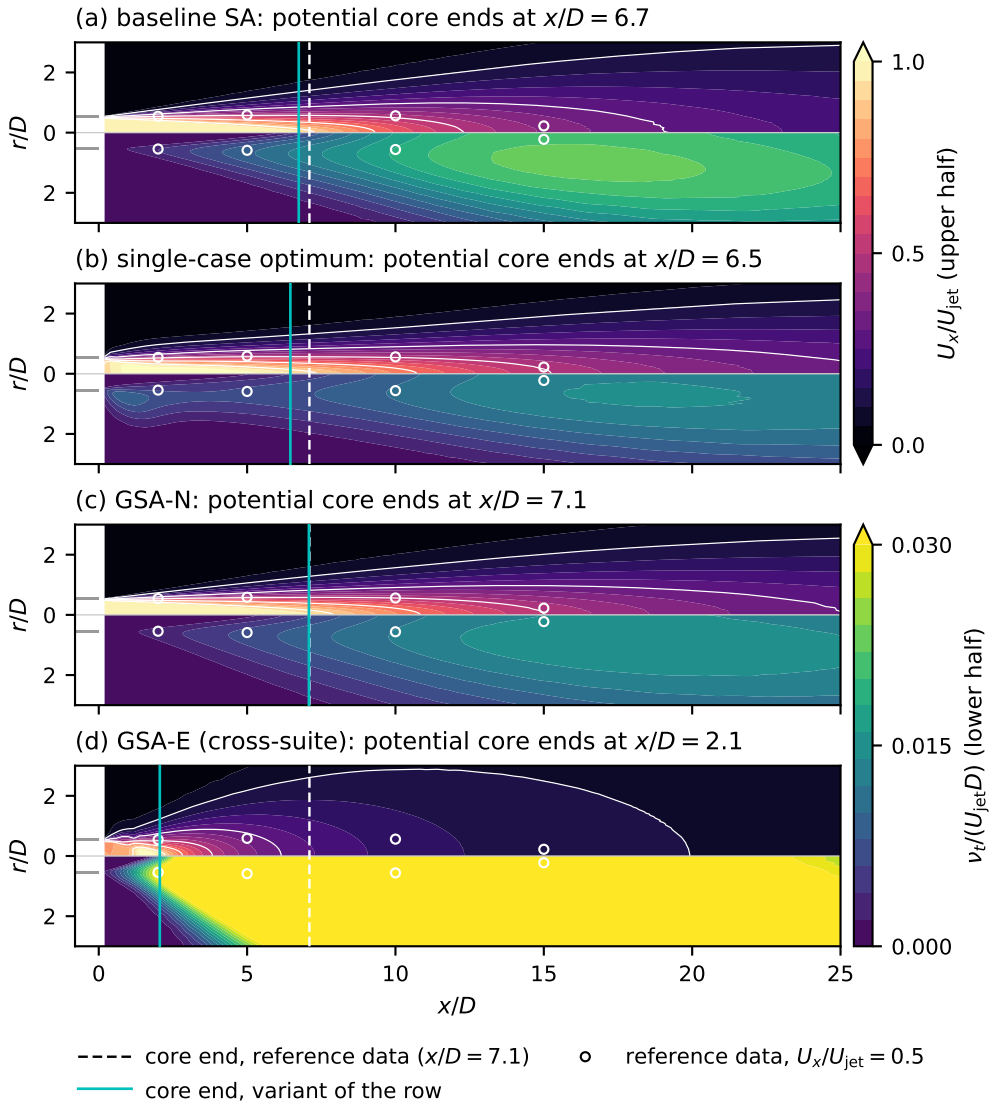

Fig. 8 ASJ in the x-r plane, one row per variant of Table 3: (a) baseline SA, (b) single-case optimum, (c) GSA-N, (d) GSA-E (cross-suite). Upper half: U_x/U_{jet} with white iso-lines 0.1-0.9; lower half (mirrored): eddy viscosity $\nu_t/(U_{\text{jet}}D)$, the colour scale saturating at 0.03. Dashed line: end of the potential core ($U_c/U_{\text{jet}} = 0.95$) in the TMR measurements (Rumsey, 2026), $x/D = 7.1$; solid line: the same criterion for the variant of the row (value in the row title); circles: measured radius of $U_x/U_{\text{jet}} = 0.5$ at $x/D = 2, 5, 10, 15$. Grey bars: nozzle lip at $r/D = 0.5$; the sampled fields start at $x/D = 0.2$. GSA-E collapses the core by $x/D \approx 2$ with an eddy viscosity several times larger than that of the other variants.

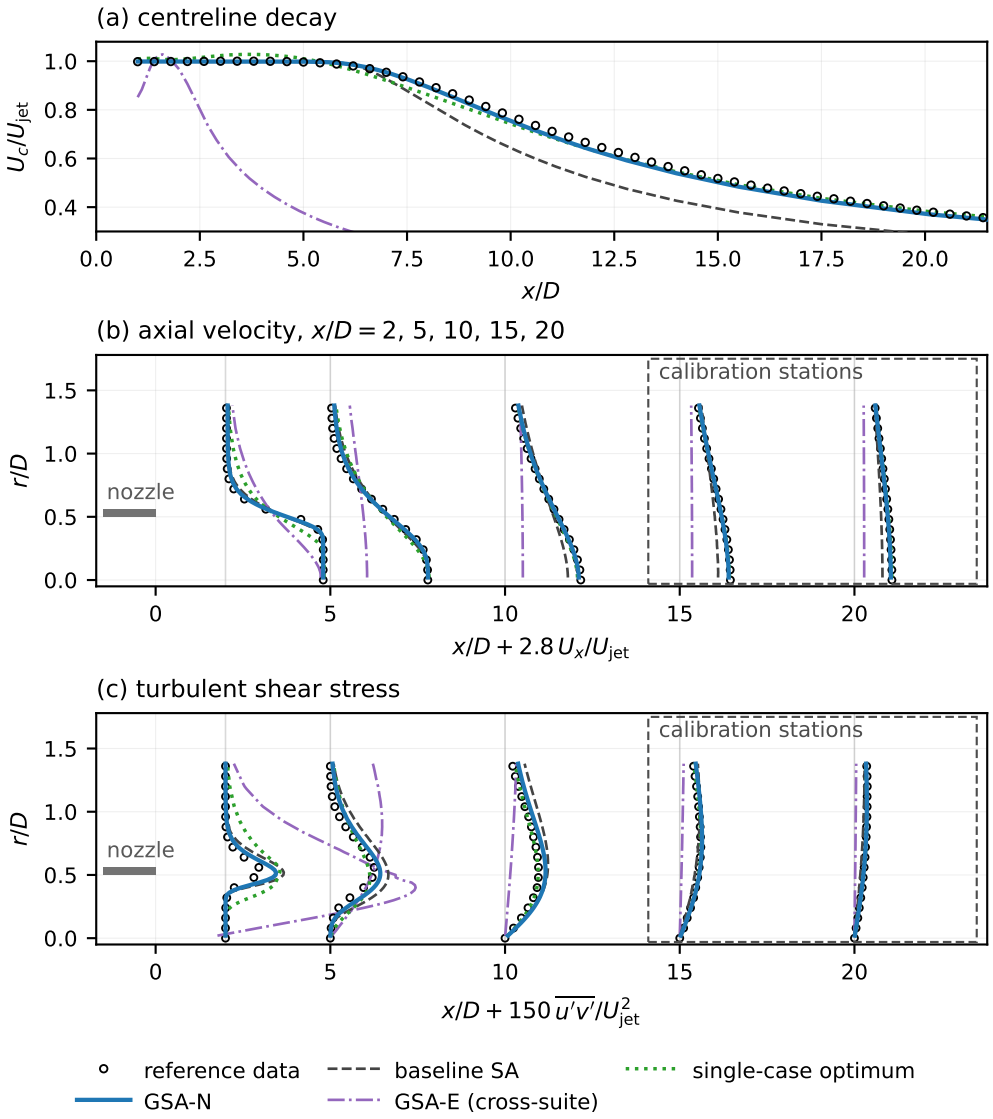

Fig. 9 ASJ: (a) centreline velocity decay against the TMR measurements (Rumsey, 2026); (b) axial-velocity and (c) turbulent shear-stress profiles at the five measured stations $x/D = 2, 5, 10, 15, 20$, each anchored at its station position (grey bar: nozzle lip at $r/D = 0.5$; dashed box: the calibration stations $x/D = 15$ and 20 that enter the cost metric, Table 1).

in the separating layer, whereas the pressure recovery and the velocity profiles also respond to its spreading. The two sets also differ in C_{rvc} (83 against 0.31), whose separate effect the present data do not isolate. The slow C_f recovery is the residual deficit of GSA-N on the hump, because the jet of the same suite keeps C_{b1O} low.

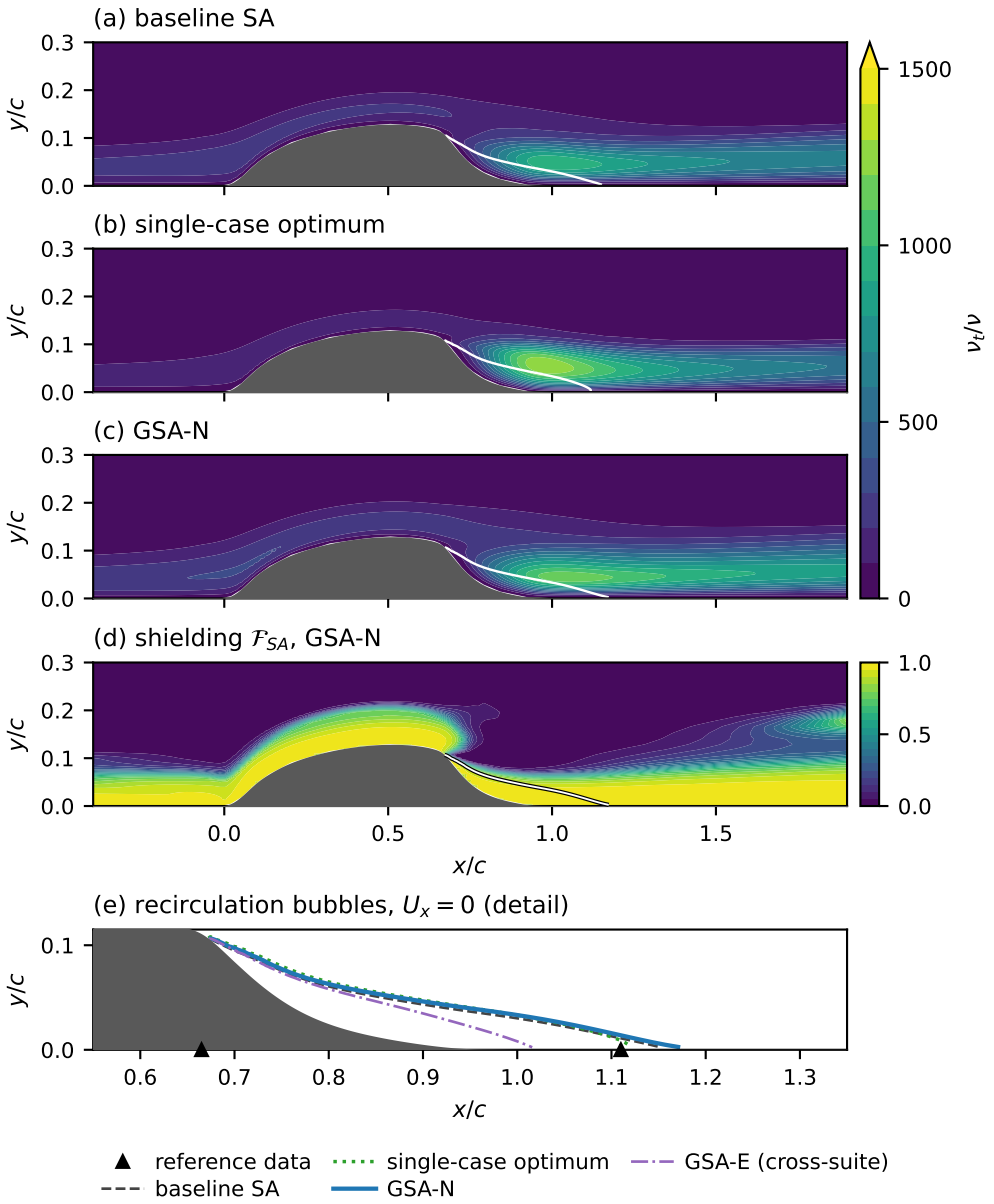

Fig. 10 2DWMH, two-dimensional fields: (a-c) eddy-viscosity ratio ν_t/ν with the $U_x = 0$ recirculation bubble (white line) for baseline SA, the single-case optimum and GSA-N; (d) shielding function of GSA-N ($\mathcal{F}_{SA} \to 0$: Outer set); (e) detail of the $U_x = 0$ bubble of every variant of Table 3, with the colours and line styles of Fig. 11; the triangles mark the separation and reattachment points measured by Greenblatt et al (2006).

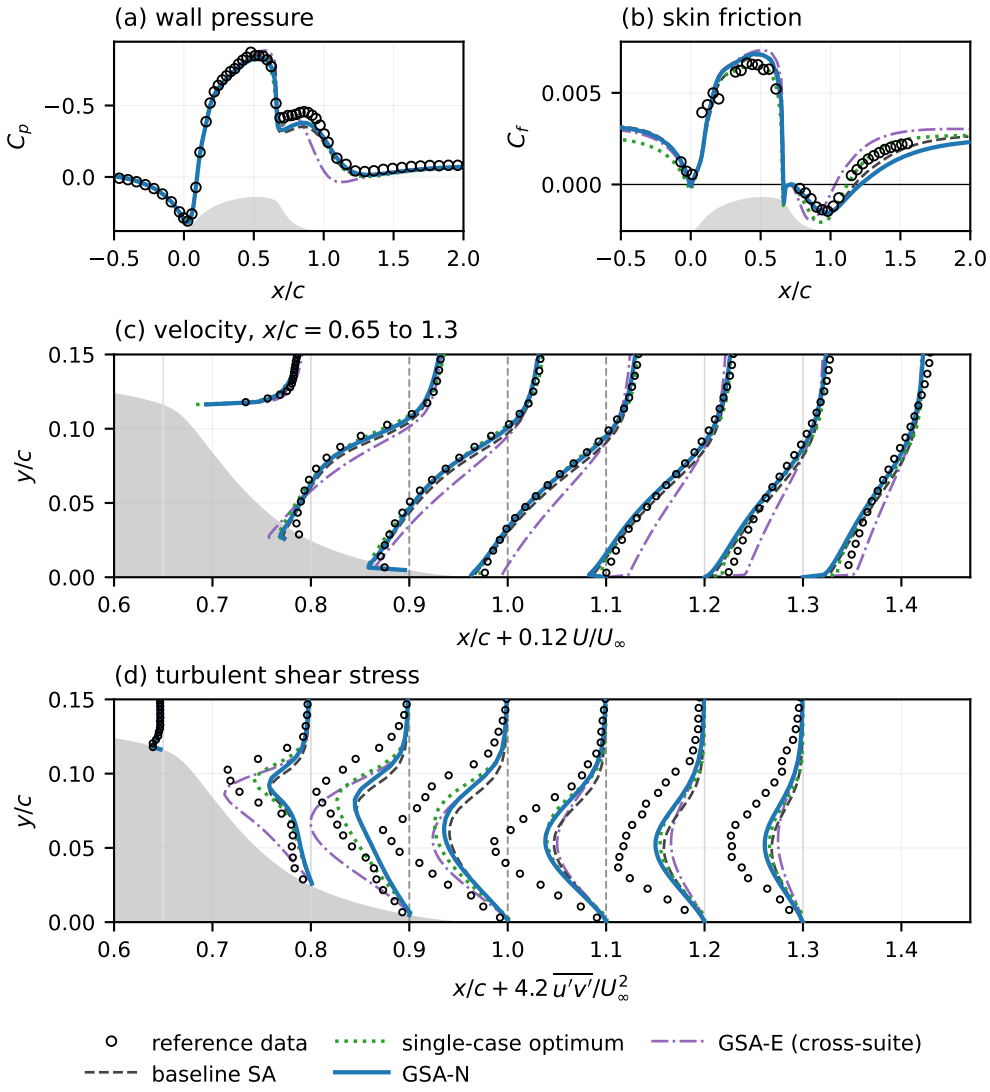

Fig. 11 2DWMH: (a) wall pressure and (b) skin friction along the wall; (c) velocity and (d) turbulent shear-stress profiles at all seven measured stations, each anchored at its station position; symbols: measurements of Greenblatt et al (2006). The profile scale ($0.12\,U/U_\infty$ and $4.2\,\overline{u'v'}/U_\infty^2$ per unit x/c) keeps each profile within about one station spacing. The dashed vertical lines mark the three stations $x/c = 0.9$, 1.0, 1.1 that enter the cost metric (Table 1); the grey silhouette is the hump surface.

4.7 ERCOFTAC suite: flow-by-flow validation

On the gentle hill ($\alpha = 1.5$) the mechanism fields of GSA-E in Fig. 12 show the shielding and the topology floor acting as designed, the floor peaking in the rotation-dominated crest-lee region and blending about 8% of the Outer set into the log layer

(Section 2.2). The profiles of Fig. 13 separate the variants clearly. Baseline SA over-separates, turning the thin, pressure-induced bubble of the DNS (Xiao et al, 2020) (at most $0.33h$ thick and essentially reattached by $x/h = 4$) into one that is too thick from $x/h = 2$ onward and still open at $x/h = 6$. GSA-E and GSA-E+QCR2024, whose bubbles nearly coincide, move markedly towards the DNS over bubble and recovery, the quadratic closure raising the shear-stress peaks (by 14% at $x/h = 2$) and lowering the cost from 0.625 to 0.584. The hill optimum (0.467) marks what a dedicated calibration reaches and is closest to the DNS shear-stress peaks. The cross-suite member GSA-N behaves like a mildly improved baseline SA, because the NASA calibration under-uses the Outer production this flow needs.

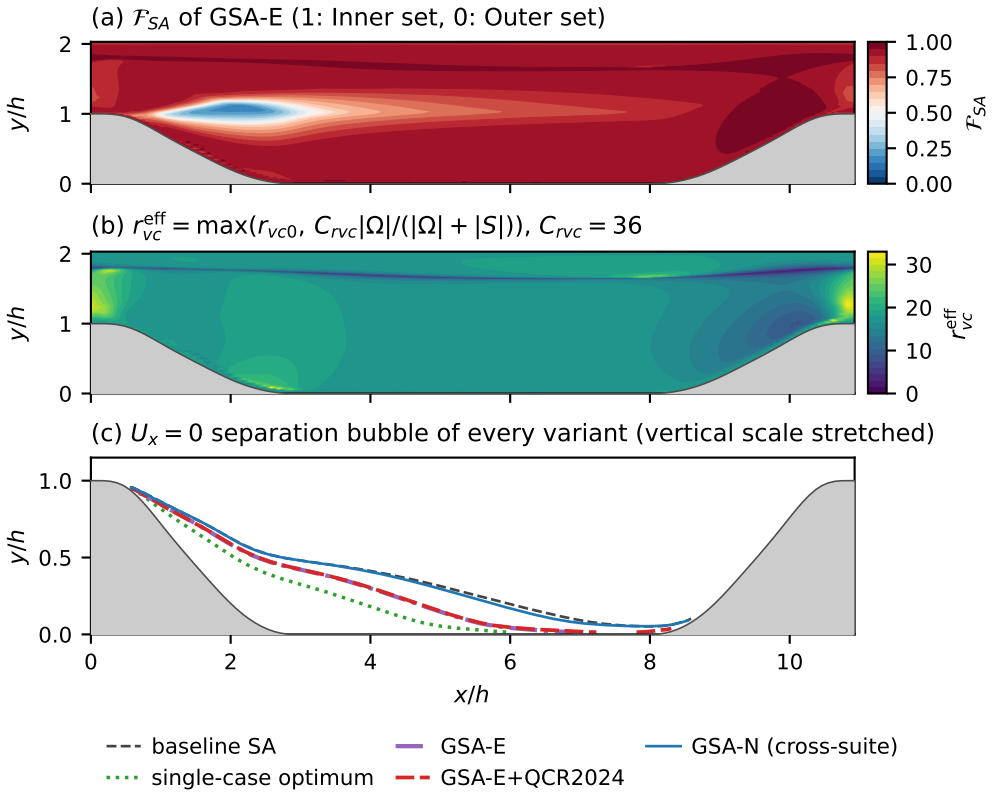

Fig. 12 Mechanism fields on the gentle parametric hill ($\alpha = 1.5$), member GSA-E: (a) shielding function ($\mathcal{F}_{SA} \to 0$, Outer set, in the separated shear layer); (b) topology-sensitive floor r_{vc}^{eff} (Eq. (13), $C_{rvc} = 36$): $\approx C_{rvc}/2$ in the sheared bulk, peaks in rotation-dominated regions, $\to r_{vc0}$ on the zero-shear core line; (c) $U_x = 0$ separation bubble of every variant of Table 3, in the colours and line styles of Fig. 13 (vertical scale stretched).

The steep hill ($\alpha = 0.5$), on which GSA-E is marginally above unity (1.07), is the harder case to improve: its separation is fixed by the steep lee slope. As shown in

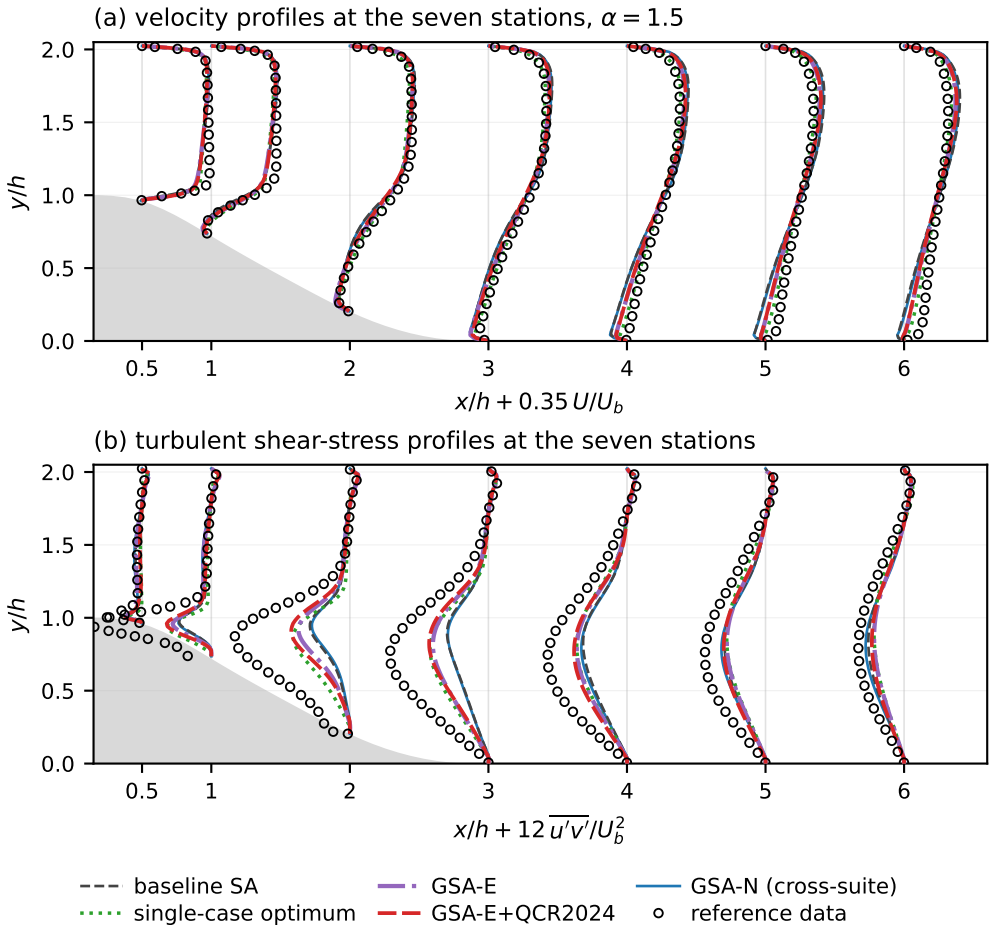

Fig. 13 Gentle parametric hill ($\alpha = 1.5$, geometry in grey): (a) streamwise-velocity and (b) turbulent shear-stress profiles at $x/h = 0.5$-6, each anchored at its station position ($0.35\,U/U_b$ and $12\,\overline{u'v'}/U_b^2$ per unit x/h), all variants against the DNS of Xiao et al (2020).

Fig. 14, the DNS bubble (Xiao et al, 2020) is large (reversed flow up to $0.8h$ thick, extending beyond $x/h = 5$), and baseline SA already reproduces it to within about $0.15h$ in thickness, with only a slightly delayed reattachment. The error left for the calibration is therefore small, and even the dedicated optimum reduces it by only 22% (0.776). GSA-E inherits the gentle-hill and curved-step reading of the Outer state and therefore slightly over-mixes the steep-hill bubble, which reattaches too early (reversed layer $0.09h$ thick at $x/h = 5$, against $0.21h$ in the DNS). GSA-E+QCR2024 keeps this bubble (reversed layer within $0.02h$ of GSA-E) but raises the shear-stress peaks towards the DNS (by up to 17%), which brings its cost below unity (0.947). The comparison also separates integral from pointwise agreement. The steep-hill optimum has the lowest seven-station cost, yet at the shear-stress peaks of the recovery region GSA-E lies as close to the DNS or closer, because the advantage of the optimum is

spread over the velocity profiles of the whole window. Resolving the two hills separately is beyond a two-state model (Section 4.1) and points to the topology sensor as the natural third handle (Section 5).

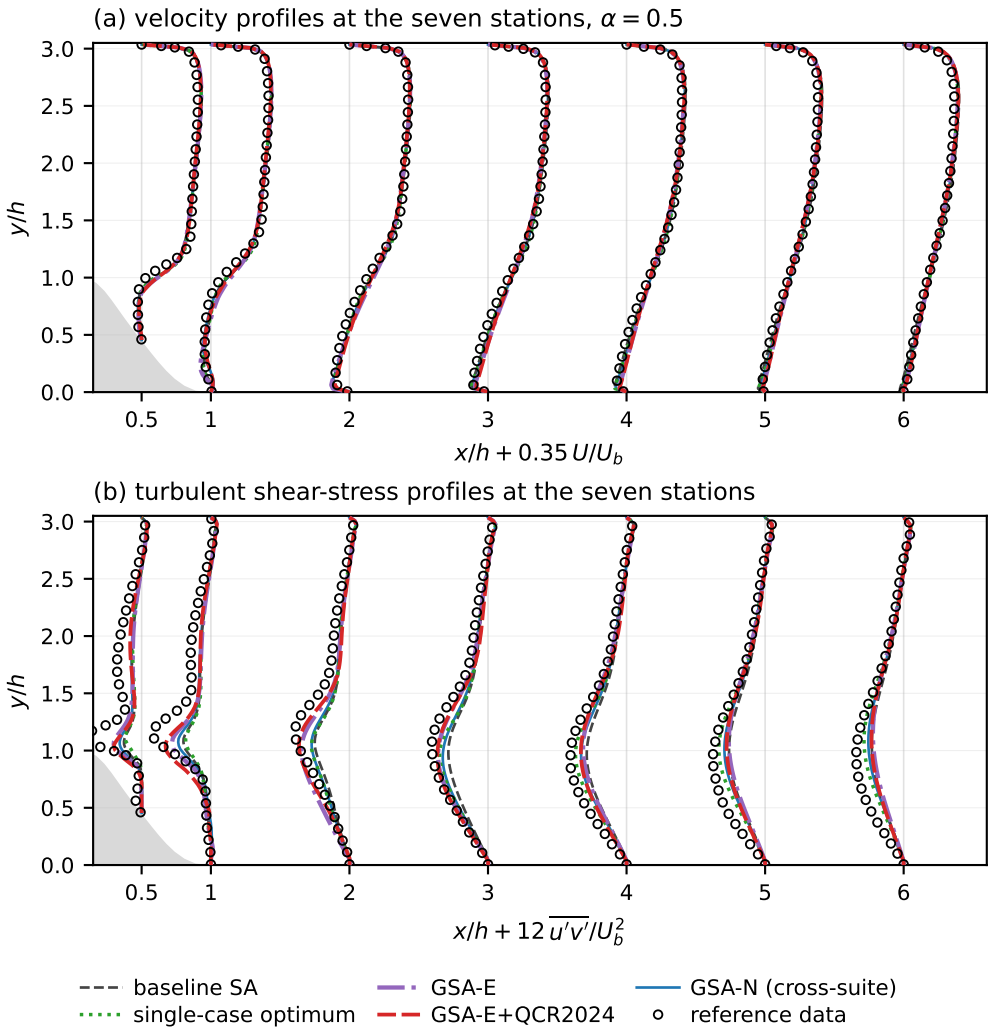

Fig. 14 Steep parametric hill ($\alpha = 0.5$, geometry in grey): (a) velocity and (b) shear-stress profiles, as in Fig. 13.

On the CBFS all calibrated variants capture the smooth-separation bubble, as shown by the C_f dip and double minimum in Fig. 15b. As shown in Fig. 15a,d, the discriminating quantities are the pressure recovery and the shear-stress peak in the detached layer, where the step optimum (0.151) comes closest to the LES of Bentaleb et al (2012), while baseline SA and GSA-E underpredict the shear-stress peak (GSA-E

by 15% and 12% at $x/H = 3.25$ and 4). GSA-E+QCR2024 overshoots the LES shear-stress peak (by up to 6%) and slows the pressure recovery (step cost 0.422, against 0.308). The step optimum is the most extreme Outer state of the suite ($C_{b1O} = 0.98$), consistent with a separated shear layer that needs the strongest mixing to reattach where the LES does. This is the easiest separated case of the suite (Table 4).

The square duct is the case that no linear variant can solve, and the one that the run4 closure was calibrated for. As shown in Fig. 16a-c, along the wall bisector every linear variant, including the duct optimum, reproduces the DNS mean-velocity profiles of Vinuesa et al (2014) to within a few per cent and the shear-stress profiles to within 15% of their peak, so the calibration differences appear only in the deviation panels d-f. None of them generates the in-plane secondary motion. As shown for GSA-E in Fig. 17a, the linear eddy-viscosity closures (baseline SA, the duct optimum and GSA-E) give machine-zero in-plane velocity ($|u_s|/U_b < 10^{-4}$, with u_s the in-plane velocity magnitude and U_b the bulk velocity), whereas GSA-E+QCR2024 develops the classical pair of counter-rotating corner vortices of Fig. 17b, with a maximum secondary velocity of 1.0% of the bulk velocity. Panels c-e of Fig. 17 follow the mechanism of GSA-E+QCR2024. The quadratic terms produce a cross-plane normal-stress difference $(R_{yy} - R_{zz})/U_b^2$ of up to 2.35×10^{-3}, identically zero for every linear closure, with the wall-parallel stress above the wall-normal one next to each wall (panel c); the cross-plane gradients of this difference drive the streamwise vorticity of the vortex pair, $|\omega_x|h/U_b$ up to 0.34 (panel d); and the resulting transport changes the streamwise velocity by -0.065 to $+0.13\,U_b$ with respect to GSA-E, accelerating the corner and slowing the near-wall region of the bisecting planes (panel e). This is Prandtl's secondary motion of the second kind, reproduced in the order in which it arises: normal-stress imbalance, streamwise vorticity, momentum redistribution. Along the wall bisector the linear variants share a $+0.02$-$0.08\,U_b$ bulge (0.11 for the cross-suite GSA-N) in the mean-velocity deviation that peaks near the wall region the corner vortices feed. As shown in Fig. 16d, QCR2024 reduces it to between -0.01 and $+0.02\,U_b$, because the secondary motion redistributes momentum towards the corners as in the DNS; this is what the duct cost measures (0.920 → 0.138).

4.8 Cross-suite validation

The cross-suite rows of Table 4 give the costs of each selected member on the other suite, plotted as the "cross-suite" variant in every flow-by-flow figure (Table 3). The two suites share no case and no reference dataset, so this is a genuine transferability test. The outcome is strongly asymmetric. GSA-N transfers well: on ERCOFTAC it improves on baseline SA on three of four cases (0.84-0.94) and degrades only the duct (1.35). GSA-E does not transfer. On NASA it preserves the wall-bounded cases (flat plate 0.015, channel 0.036), which the Inner state and the C_{w1} reconstruction protect by construction, but fails on the two cases governed by the Outer state, the jet (14.1) and the hump (2.6). The most likely cause is the Outer-production conflict of Section 4.1 (Table 5). With $C_{b1O} = 0.44$, four times that of GSA-N, the jet potential core collapses by $x/D \approx 2$ (Figs. 8d and 9a). On the ERCOFTAC suite only the steep hill favours a low C_{b1O}, so the mean-cost selection lets it grow, whereas on the NASA

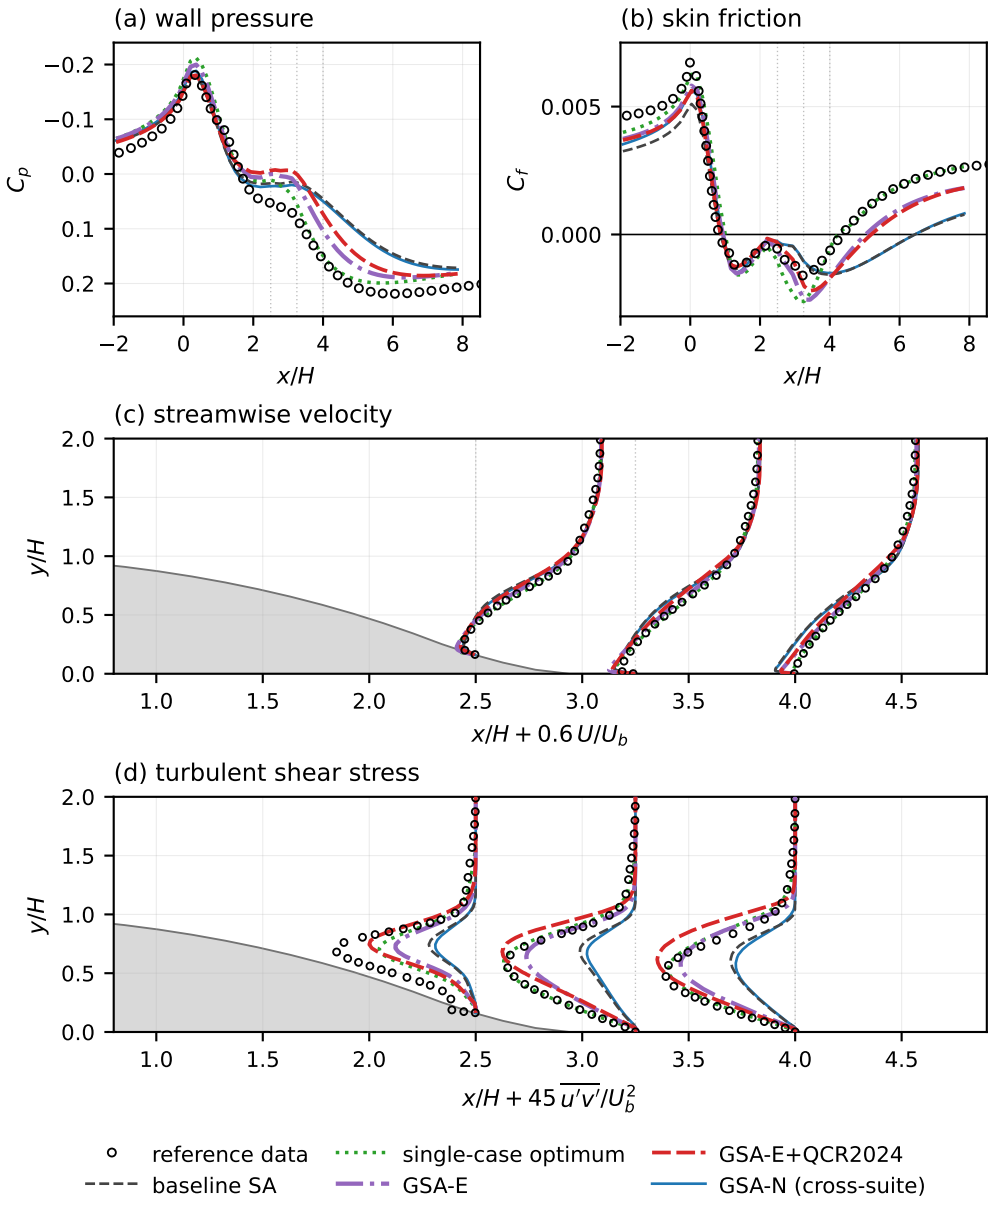

Fig. 15 CBFS: (a) wall pressure coefficient and (b) skin-friction coefficient along the step wall, with the three profile stations marked by dotted lines; (c) streamwise-velocity and (d) turbulent shear-stress profiles at $x/H = 2.5$, 3.25 and 4, anchored at their station positions (scales $0.6\,U/U_b$ and $45\,\overline{u'v'}/U_b^2$ per unit x/H), all variants against the LES of Bentaleb et al (2012); step geometry (grey) to scale, reference markers thinned.

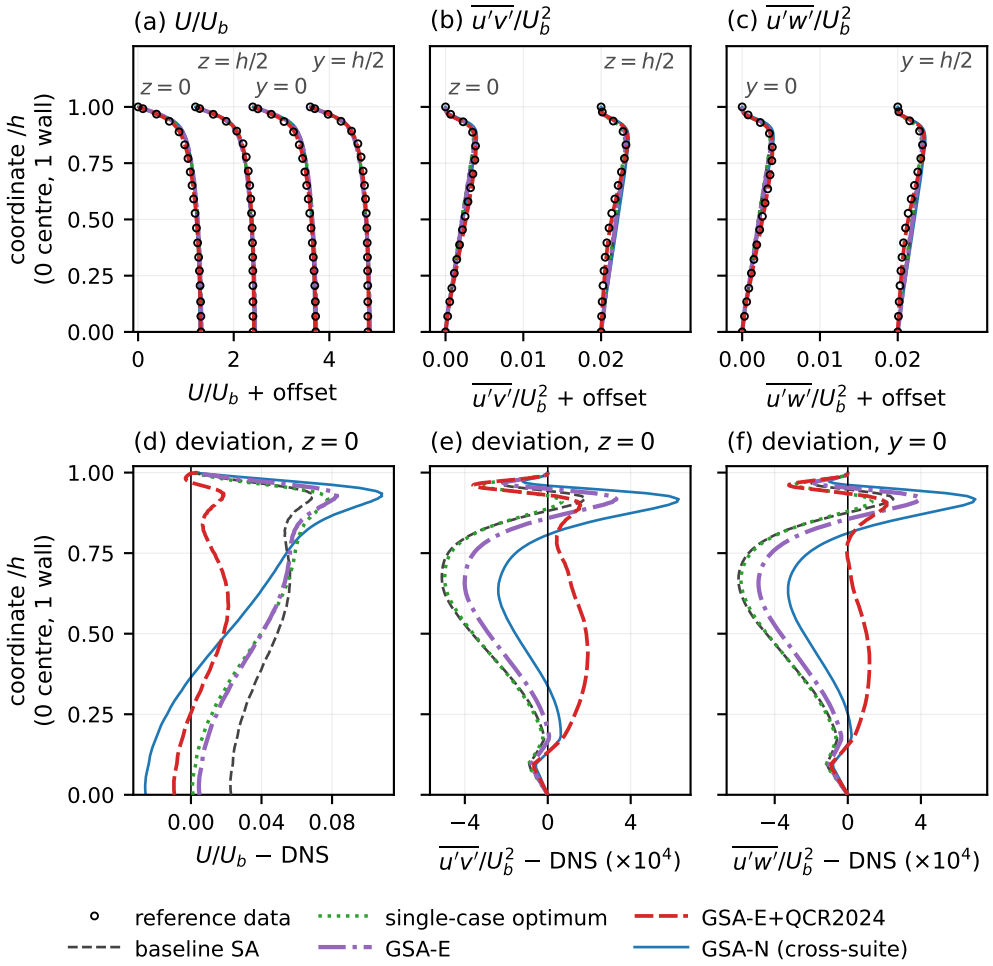

Fig. 16 Square duct ($\mathrm{Re}_\tau = 180$): (a-c) mean velocity on the four lines of the calibration metric ($z = 0$, $z = h/2$, $y = 0$, $y = h/2$), primary shear stress on $z = 0$ and $z = h/2$ and secondary shear stress on $y = 0$ and $y = h/2$, against the DNS of Vinuesa et al (2014); the coordinate along each line is on the vertical axis (0 centre, 1 wall) and the lines are shifted horizontally by 1.2 (U/U_b) and 0.02 (stresses) for readability. (d-f) pointwise deviation from the DNS on the first line of each panel (unshifted; stresses in units of $10^{-4}\,U_b^2$).

suite the jet bounds it from above and the hump constrains separation, which is why GSA-N remains reasonable on the hills and step.

4.9 Validation on configurations unseen in calibration

As a final, stricter test, the two selected members are applied with frozen coefficients to three configurations that no campaign saw: two external flows at conditions outside the calibration envelope and one fully three-dimensional junction flow. Baseline SA

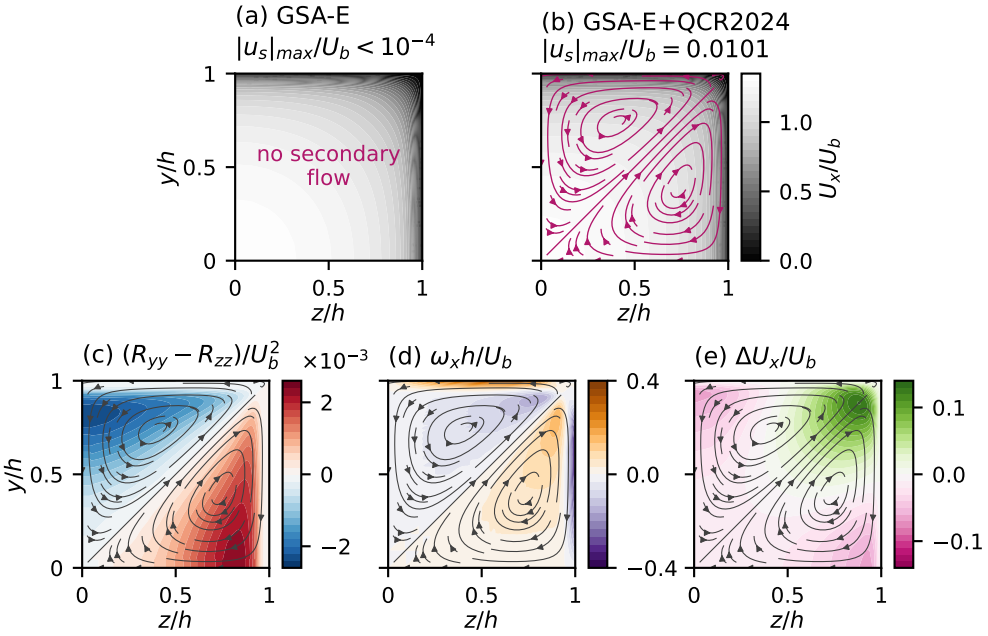

Fig. 17 Square duct ($\mathrm{Re}_\tau = 180$, quarter section, walls at $y/h = z/h = 1$). (a-b) Streamwise velocity U_x/U_b (greyscale, common scale) and in-plane streamlines for (a) GSA-E, without secondary flow (baseline SA and the duct optimum are indistinguishable and not shown), and (b) GSA-E+QCR2024, with the corner-vortex pair. (c-e) Mechanism of the quadratic term for GSA-E+QCR2024: normal-stress difference $(R_{yy} - R_{zz})/U_b^2$ of the modelled Reynolds stress (zero for the linear closure), streamwise vorticity $\omega_x h/U_b$ and change of streamwise velocity with respect to GSA-E, $\Delta U_x/U_b$, with the secondary-flow streamlines.

is run through the same implementation (Inner = Outer = baseline, $C_{rvc} = 0$), and GSA-E also with the calibrated QCR2024 closure (GSA-E+QCR2024).

The NACA 0012 polar (TMR case 2DN00) is computed through stall at $\mathrm{Re}_c = 6 \times 10^6$ and $\alpha = 0, 10, 15, 16, 17, 18, 19°$ on the 897×257 C-grid of the TMR (Rumsey, 2026) (229k cells, far field at $\sim 500c$, $y^+_{\max} < 0.6$). Figure 18 compares the lift and drag of the four variants with the tripped measurements of Ladson (1988). Each incidence is iterated until the force coefficients are stationary (half-window lift shift over the last 5000 iterations below 3.5×10^{-3}, i.e. 0.2% of C_L), which requires up to 1.4×10^5 pseudo-time steps near and beyond stall, because incomplete convergence overstates post-stall lift by up to 0.1 and was found to hide the stall onset entirely. Baseline SA reproduces the reference SA solutions of the TMR (Rumsey, 2026) at $\alpha = 10°$ ($C_L = 1.0778$, $C_D = 0.01212$) and $\alpha = 15°$ ($C_L = 1.5461$) to 0.1%, validating the numerical set-up. In the linear range every variant lies on the data. At $\alpha = 10°$ the lift coefficients are 1.079 (baseline SA), 1.081 (GSA-N), 1.093 (GSA-E) and 1.081 (GSA-E+QCR2024) against the measured ≈ 1.08 (Ladson, 1988), and the zero-incidence drag spans 0.0083-0.0086 against the measured 0.0081. The generalised calibration thus preserves the SA behaviour on an unseen attached external flow. Around stall the

members separate along the coefficient physics of Sections 4.1 and 4.8. The experiment stalls at $C_{L,\max} = 1.61$-1.63 ($\alpha \approx 17°$) (Ladson, 1988); baseline SA stalls at $C_{L,\max} = 1.73$ (18°) and GSA-N at 1.74 (18°), both with a lift drop at 19°. The NASA member thus inherits the stall onset of baseline SA essentially unchanged, one degree late and 6-8% high, the documented accuracy limit of fully turbulent steady two-dimensional RANS on a case whose wind-tunnel stall is itself affected by three-dimensional side-wall effects (McCroskey, 1987). GSA-E does not stall within the computed range: it reaches $C_L = 1.90$ at 19° (an upper bound, Fig. 18b), most likely because its large Outer production ($C_{b1O} = 0.44$) delays the growth of the trailing-edge separation that triggers stall. The quadratic stress restores the stall: GSA-E+QCR2024 reaches $C_{L,\max} = 1.74$ at 18° and drops to 1.68 at 19°, like baseline SA, with the largest drag of all variants from 15° on (0.040 at 18°, against 0.034 for GSA-E), which suggests, without field data to confirm it, a larger trailing-edge separation. In attached flow it returns GSA-E to the lift of baseline SA and lowers its zero-incidence drag by 2%. The GSA-E+QCR2024 points use the stabilised formulation of Section 2.3 ($\alpha_{\mathrm{QCR}} = 0.05$, $\lambda_{\mathrm{QCR}} = 1$). Its limiter acts only mildly: it clips 2-12% of the turbulent cells, mostly in the Outer region, with a median scale factor of 0.99-1.00, so the converged states are essentially the QCR2024 closure.

The axisymmetric transonic bump (ATB) of Bachalo and Johnson (1986) at Mach number $M = 0.875$ and $\mathrm{Re}_c = 2.763 \times 10^6$ combines a supersonic pocket, shock-induced separation and reattachment, and compressibility physics absent from the calibration envelope. The case is run with the compressible variant of the model coupled to the density-based HiSA solver (Heyns et al, 2014) (advection upstream splitting method AUSM$^+$-up with total-variation-diminishing reconstruction, characteristic far-field conditions, and pseudo-time marching with generalised minimal residual iterations and a lower-upper symmetric Gauss-Seidel preconditioner) on an axisymmetric wedge mesh. Because HiSA assembles its viscous fluxes independently of the stress-divergence interface of the turbulence model, the QCR2024 stress is injected explicitly into the momentum and energy residuals, a small solver-side coupling released with the model. The three linear variants converge the pseudo-time residuals by three orders of magnitude without special treatment; GSA-E+QCR2024 is warm-started from the converged GSA-E state and uses the same stabilised formulation. As shown in Fig. 19a, the computed wall pressure reproduces the suction plateau measured by Bachalo and Johnson (1986), with peak $C_p = -0.83$ (baseline SA), -0.85 (GSA-N), -0.84 (GSA-E) and -0.82 (GSA-E+QCR2024) against the measured ≈ -0.8, and the shock, located at the steepest pressure rise, at $x/c = 0.68$, 0.70, 0.68 and 0.67 against the measured 0.66. This is the canonical, slightly aft-shifted shock of the SA family on this case, which the two blending members barely move, whereas the quadratic closure moves it about $0.01c$ upstream, to within $0.01c$ of the measurement. As shown in Fig. 19b,c, the velocity and turbulent shear-stress profiles across the interaction follow the measurements at all six stations, with the model-to-model spread confined to the $\overline{u'w'}$ peak in the separated region, which GSA-E+QCR2024 predicts closest to the measurements downstream of the shock. The blending members therefore do not degrade the shock/boundary-layer-interaction physics that the calibration never saw.

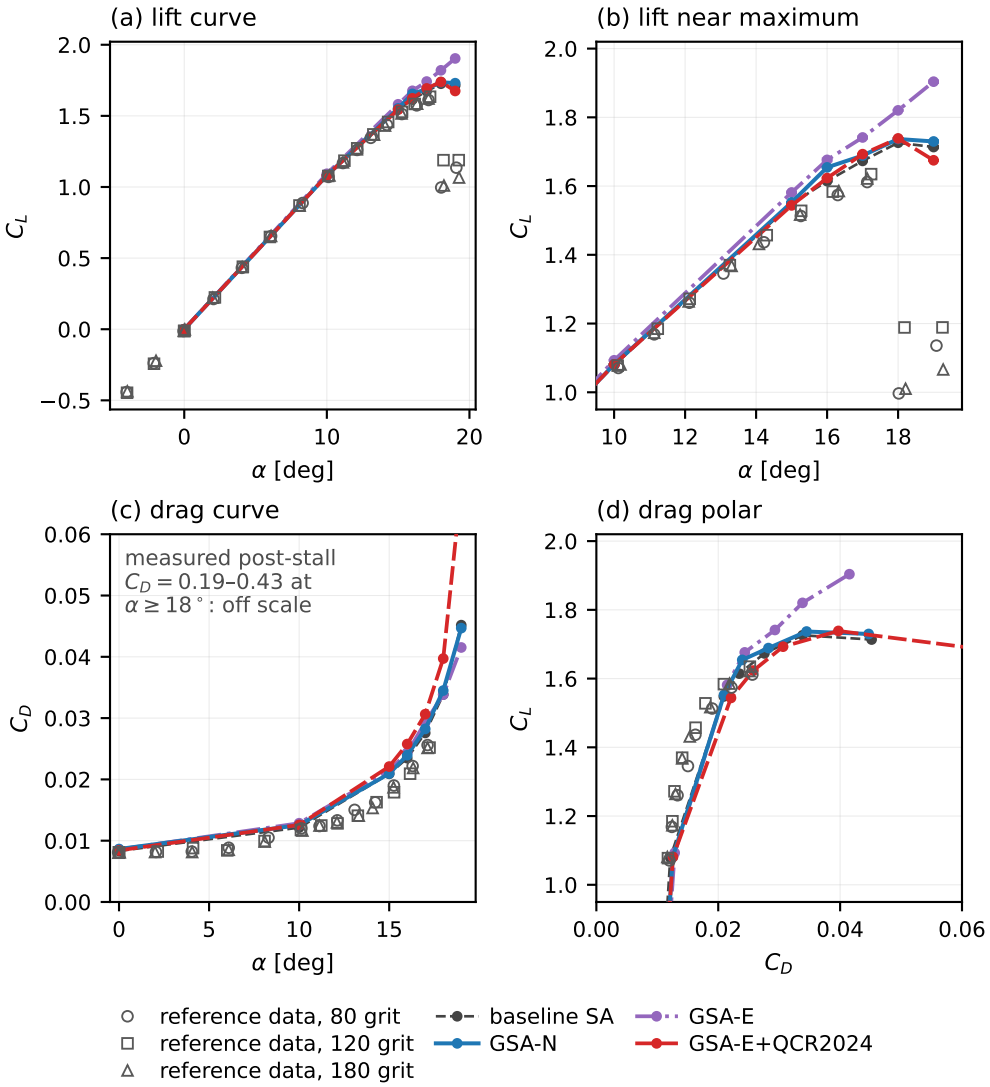

Fig. 18 2DN00 at $\mathrm{Re}_c = 6 \times 10^6$: (a) lift curve, (b) lift near maximum, (c) drag curve, (d) drag polar; the four model variants vs the tripped measurements of Ladson (1988) (grey symbols, three trip grits). The measured post-stall drag (C_D = 0.19-0.43 at $\alpha \geq 18°$) and the 19° drag of GSA-E+QCR2024 (0.068) are off scale in (c) and (d). Error bars in (b): residual-transient band (half the lift envelope over the last 5000 iterations), visible only where a slow limit cycle persists. The 19° point of GSA-E is an upper bound: its lift still relaxes slowly downward at the end of the iterations (Section 4.9).

The wing-body junction (WBJ) of Devenport and Simpson (1990) ($\mathrm{Re}_T = 1.15 \times 10^5$ based on the maximum thickness T of the 3:2 elliptic-nose/NACA 0020 wing) is the only fully three-dimensional test, and one of the configurations of the ERCOFTAC

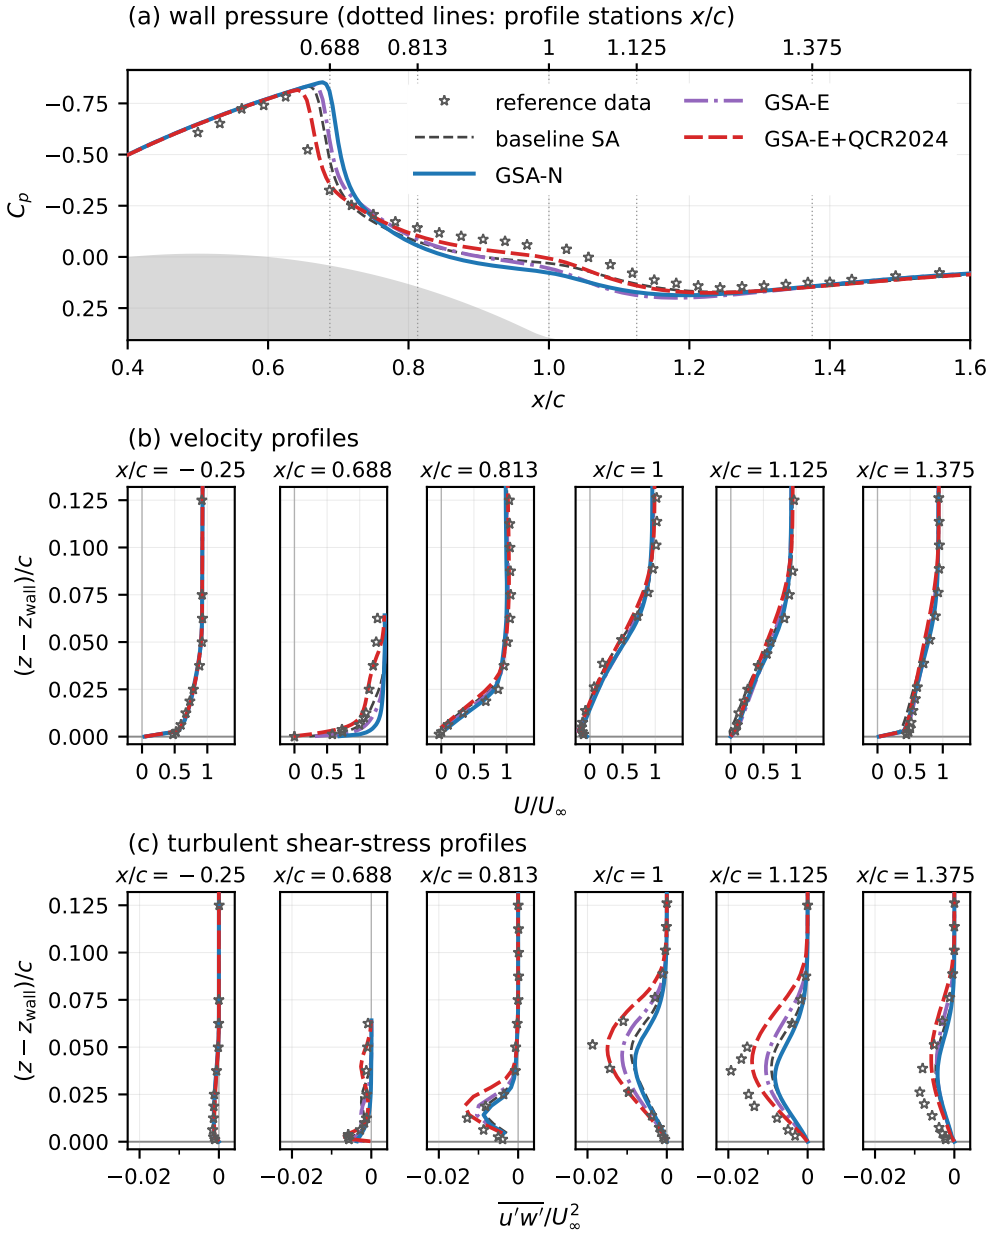

Fig. 19 ATB at $M = 0.875$, $\mathrm{Re}_c = 2.763\times10^6$, computed with the density-based HiSA solver (Heyns et al, 2014), against the measurements of Bachalo and Johnson (1986): (a) wall pressure over the interaction, with the bump in grey and the profile stations marked by dotted lines and labelled at the top (the upstream reference station $x/c = -0.25$ lies outside the panel); (b) velocity and (c) turbulent shear-stress profiles at the six measured stations, each in its own panel, against the distance from the local wall.

data-driven turbulence-modelling challenge (Buchanan et al, 2026). It is computed on an 8.6M-cell mesh and compared, in Fig. 20, with the velocity and Reynolds-stress profiles measured by Devenport and Simpson (1990), normalised by the reference velocity U_{ref} of the experiment, at ten stations along the upstream symmetry plane, $x/T = -0.45$ to -0.01, approaching the wing leading edge. Averaged over the ten stations within $y/T \leq 0.2$, GSA-N reduces the rms error with respect to the measurements by 15% in streamwise velocity and 18% in the wall-normal component relative to baseline SA, and GSA-E by 8% and 10%; the shear-stress error is marginally larger for the calibrated members (+7% on a small absolute level). GSA-E+QCR2024 reduces the velocity errors by 9% and 19%, the largest reduction in the wall-normal component, but raises the shear-stress error by 24%, because it overpredicts the small measured stress at $x/T = -0.10$ and -0.05. The corner region, the design domain of the quadratic closure, develops downstream of the junction and is not covered by these upstream stations.

Figure 21 adds the symmetry-plane view of the horseshoe vortex that these profiles sample. All four closures produce the classical single-vortex topology. The plate boundary layer separates at $x_s/T = -0.45$ (baseline SA and GSA-N) and -0.46 (GSA-E and GSA-E+QCR2024), rolls into a primary vortex whose centre, located with the Γ_1 swirl criterion (Graftieaux et al, 2001), lies at $x/T = -0.25$, -0.23, -0.24 and -0.24 in the same order, 0.03-0.04 T above the plate, and feeds a thin layer of reversed flow that reaches the junction corner, where the near-wall flow turns forward again at $x/T \approx -0.02$. The measurements of Devenport and Simpson (1990) bracket this picture: reversed flow at the wall appears between $x/T = -0.35$ and -0.30, peaks at $-0.52\,U_{\mathrm{ref}}$ at $x/T = -0.15$ inside a layer $0.04\,T$ thick, and the wall-normal velocity at $y/T = 0.05$ changes sign between $x/T = -0.25$ and -0.20, which places the measured vortex centre near $x/T \approx -0.2$. Every closure therefore predicts a vortex that is slightly larger and sits further upstream than the measured one. GSA-N brings the centre closest to the measurements, but the topology is unchanged, so the error reduction quantified above is a gain in the velocity distribution around an essentially common vortex, not a different vortex system. The remaining offset is not a matter of coefficients. The calibrated sets move the vortex centre by at most $0.02\,T$, less than the residual offset of about $0.03\,T$, and the QCR2024 anisotropy leaves the separation point unchanged and moves the vortex centre by less than $0.01\,T$ towards the data. The measured vortex is dominated by a low-frequency bimodal unsteadiness (Devenport and Simpson, 1990), which a steady closure, linear or quadratic, represents only through its time-averaged effect; the discrepancy is therefore most likely a model-form limitation of steady RANS for this flow rather than a calibration deficit.

The blending members improve these three flows at most modestly, consistent with remaining errors that arise in regimes that the calibration did not contain. The attached boundary layers of the aerofoil and of the bump lie in the Inner state, whose log-layer balance the C_{w1} reconstruction protects (Section 2.1), so both members stay close to baseline SA there (lift at $\alpha = 10^\circ$ within 1.3%). The remaining errors are set by mechanisms absent from both suites: the trailing-edge separation that triggers stall, whose onset changes with the Outer production and with the constitutive relation; the shock and its induced separation; and the unsteady horseshoe vortex. GSA-N

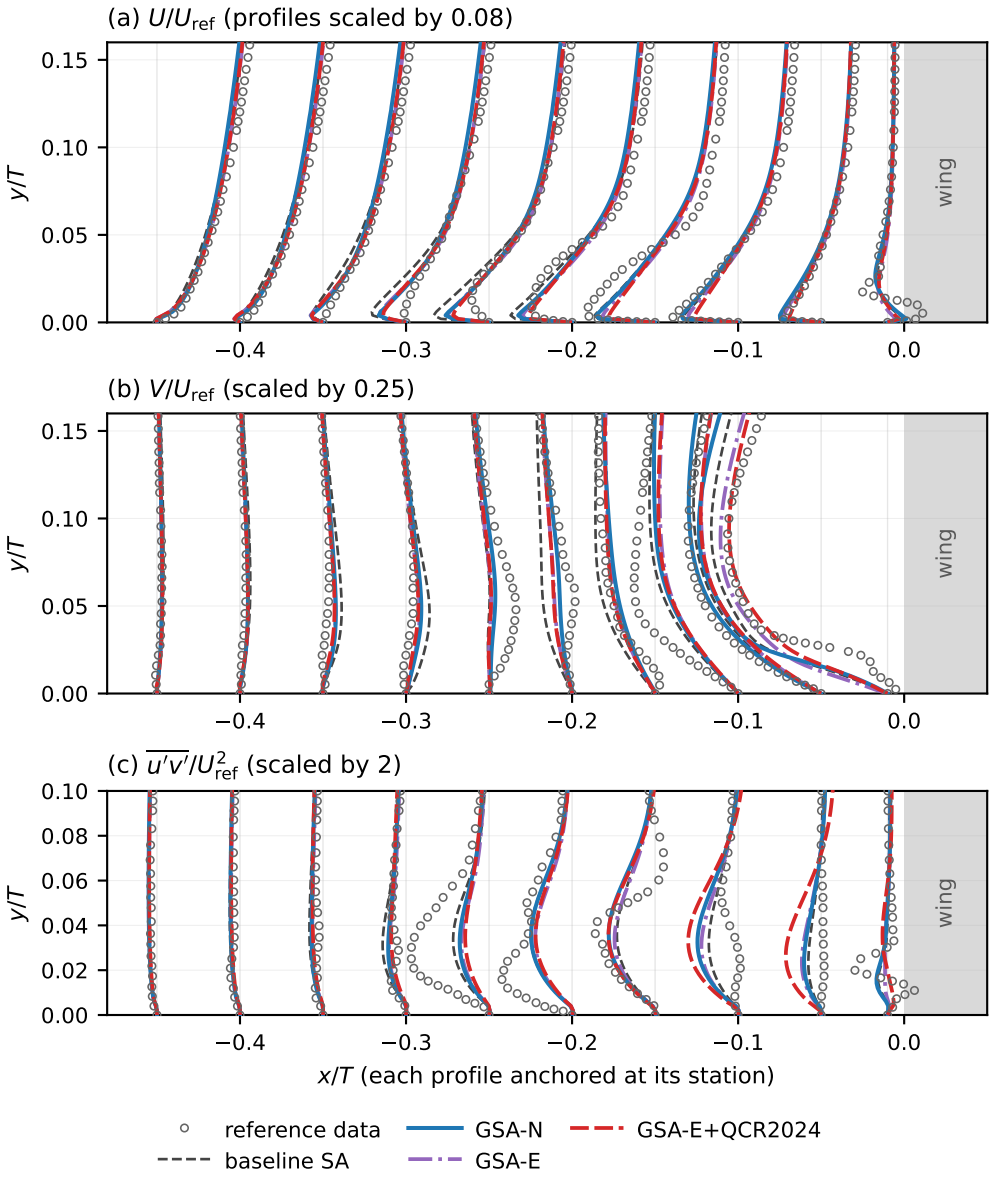

Fig. 20 WBJ (Devenport and Simpson, 1990): measured vs computed profiles of (a) streamwise velocity, (b) wall-normal velocity and (c) turbulent shear stress along the upstream symmetry plane approaching the wing leading edge (grey band), $x/T = -0.45$ to -0.01. Each profile is anchored at its station and scaled by the factor given in the panel title; the vertical range is limited to the region where the closures differ ($y/T \leq 0.16$ in (a) and (b), $y/T \leq 0.10$ in (c)), and the measurement markers are thinned for legibility.

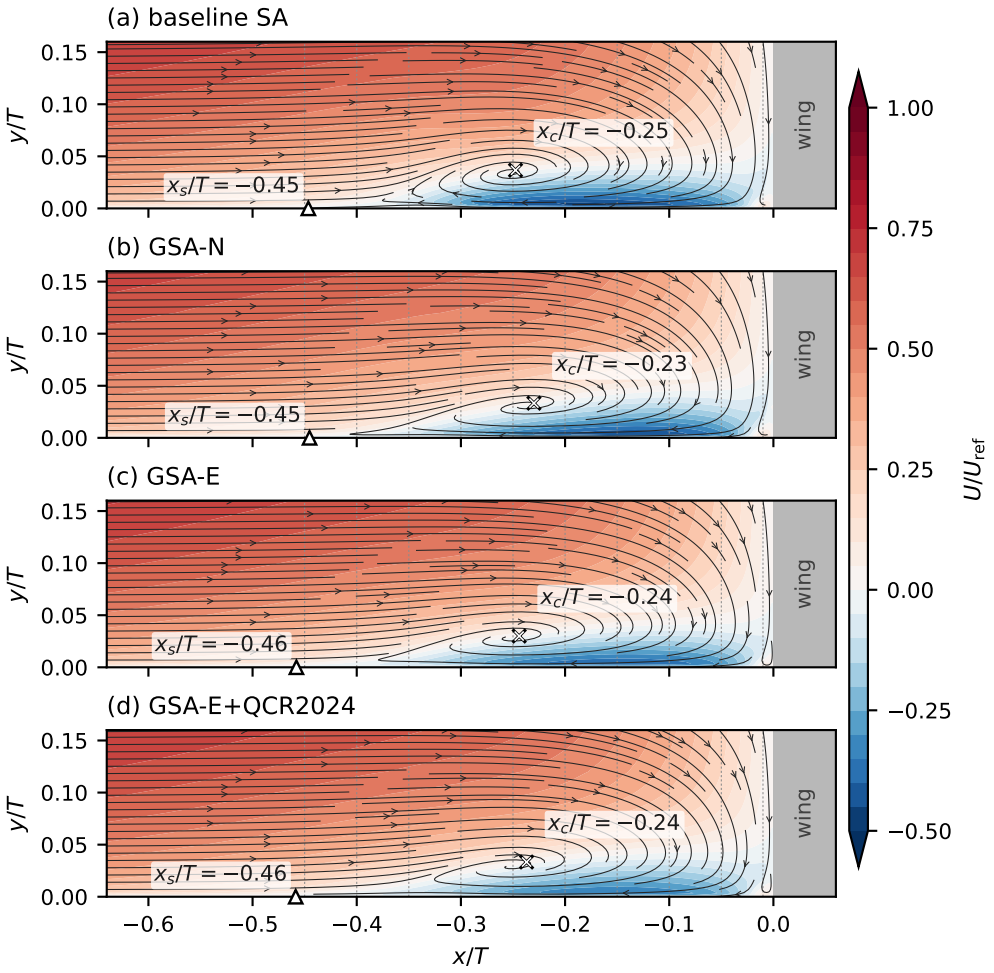

Fig. 21 WBJ: symmetry-plane view of the horseshoe vortex ahead of the wing nose (grey) for (a) baseline SA, (b) GSA-N, (c) GSA-E and (d) GSA-E+QCR2024: streamwise velocity U/U_{ref} (blue: reversed flow) and in-plane streamlines. Triangle: separation point x_s of the plate boundary layer (first $U = 0$ crossing on the lowest grid level); cross: vortex centre from the Γ_1 swirl criterion (Graftieaux et al, 2001); dotted lines: the ten measurement stations of Fig. 20.

and GSA-E+QCR2024 both preserve the stall of baseline SA, GSA-E+QCR2024 also improves the shock position and GSA-N gives the most balanced junction profiles. Single-case calibrations on these flows, as in Section 4.1, would identify the coefficient states they require and are left to future work.

4.10 Realisability of the modelled stress: anisotropy-invariant analysis

We finally check that the Reynolds stress modelled by the quadratic closure remains physically admissible, both in the square duct, where it generates the secondary flow, and on the three-dimensional junction. A Reynolds-stress tensor is realisable if it is positive semi-definite (Schumann, 1977). In terms of the anisotropy tensor $b_{ij} = \mathrm{dev}(R_{ij})/(2k)$, this holds if and only if every eigenvalue of b_{ij} lies in $[-1/3, 2/3]$; this eigenvalue test gives all the counts below. The states are displayed in the anisotropy-invariant map of Lumley (1979), in the coordinates $\eta = (b_{ij}b_{ji}/6)^{1/2}$ and $\xi = (b_{ij}b_{jk}b_{ki}/6)^{1/3}$. Realisable states fill the curved triangle bounded by the axisymmetric-contraction ($\eta = -\xi$) and axisymmetric-expansion ($\eta = \xi$) limits, which meet at the isotropic point, and by the two-component curve $\eta = (1/27 + 2\xi^3)^{1/2}$, which ends at the one-component vertex $(1/3, 1/3)$.

One-equation models carry no turbulent kinetic energy (TKE), so k must come from a proxy. We take it from the isotropic term of the QCR2024 relation, Eq. (16), with the published $C_{cr3} = 2.15$ (Tamaki and Kawai, 2024): $k = \frac{3}{2}C_{cr3}\nu_t|\Omega|$, the same for all variants; for GSA-E+QCR2024 it is exactly half the trace of the modelled stress. Every term of the modelled stress, linear and quadratic, is proportional to ν_t, so the eddy viscosity cancels in b_{ij}: the diagnostic tests the algebraic map from the mean velocity gradient to the anisotropy implied by each closure, evaluated on its own solution, and the realisability test applies to the stress $\frac{2}{3}k\delta_{ij} + \mathrm{dev}(R_{ij})$ built with this proxy. For any linear eddy-viscosity closure, $b_{ij} = -\nu_t S_{ij}/k = -S_{ij}/(\frac{3}{2}C_{cr3}|\Omega|)$, so the state depends only on the shape of the local velocity gradient: $\eta = \eta^*|S|/|\Omega|$, with $\eta^* = (3\sqrt{3}\,C_{cr3})^{-1} = 0.090$, and ξ has the sign of $-\det(S_{ij})$. A linear closure therefore approaches the isotropic point only where rotation dominates, $|\Omega| \gg |S|$. In a parallel shear flow, $|S| = |\Omega|$ and $\det(S_{ij}) = 0$: every linear closure sits at the pure-shear point $(\xi, \eta) = (0, \eta^*)$, with eigenvalues $\pm(3C_{cr3})^{-1} = \pm 0.155$ and 0. The height η^* scales with $1/C_{cr3}$, so it is set by the proxy, not by the model; only departures from the pure-shear point carry information on the closure.

As shown in Fig. 22a,b, every cell of the square-duct solutions of baseline SA and GSA-E falls on the pure-shear point, because the mean flow is unidirectional; this holds independently of position, coefficients and blending. The collapse is the invariant-map signature of the structural deficit of Sections 4.3 and 4.5: a linear closure admits a single anisotropy state in a parallel flow, with equal cross-plane normal stresses, and therefore no mechanism to drive Prandtl's secondary motion of the second kind. As shown in Fig. 22c, GSA-E+QCR2024 breaks the degeneracy. In a parallel shear flow with velocity gradient G, its quadratic terms add the normal stresses $\nu_t G\,(2C_{cr1} + C_{cr2}/12,\ -2C_{cr1} + C_{cr2}/12,\ -C_{cr2}/6)$ in the streamwise, wall-normal and spanwise directions, which with the constants of GSA-E+QCR2024 move the pure-shear state to $(\xi, \eta) = (0.094, 0.122)$, with eigenvalues $(-0.190, -0.039, 0.228)$, on the axisymmetric-expansion side of the map. The duct states cluster around this point ($\xi = 0.088$-0.103 and $\eta = 0.120$-0.125 between the 1st and 99th percentiles). The streamwise normal stress now dominates (median $b_{xx} = 0.159$), as in wall turbulence. Of the two cross-plane stresses, the wall-parallel one exceeds the wall-normal one by $\nu_t G\,(2C_{cr1} - C_{cr2}/4)$, positive for GSA-E+QCR2024 and in every duct cell off the corner bisector:

this is the imbalance that drives the corner vortices (Fig. 17c). The strain term alone would reverse it (Section 4.5). The intermediate (spanwise) eigenvalue stays negative because the streamwise stress takes the largest share, which places the states on the expansion side. The smallest eigenvalue reaches at most -0.195, about a fifth of the way from the pure-shear value of the linear closures to the two-component limit $-1/3$, the largest at most 0.242, and the minimum vertical distance to the two-component curve is 0.072. No cell leaves the triangle, including the centre cell, where the velocity gradient almost vanishes by symmetry: the modelled stress is realisable in every cell outside the wall-adjacent layer. As shown in Fig. 22d, the excess anisotropy $\eta - \eta^*$ lies between 0.028 and 0.037: it is smallest on the corner bisector in the duct core, largest in the core next to the symmetry planes, and close to its pure-shear value 0.033 in the near-wall layers.

The duct is two-dimensional in the mean, so the analysis is repeated on the fully three-dimensional WBJ of Section 4.9. The modelled stress is sampled at ∼88k points of a lattice around the junction ($-0.5 \le x/T \le 4.5$, $y/T \le 0.5$, $|z|/T \le 1.5$) for the converged solutions of baseline SA, GSA-E and GSA-E+QCR2024, whose anisotropy-invariant maps are shown in Fig. 23.

In three dimensions the strain and the vorticity no longer balance, and the proxy reaches its limit: where $|S| > |\Omega|$ the deviatoric stress $-2\nu_t S_{ij}$ grows relative to k, and in irrotational strain k vanishes while the stress does not. For a linear closure the eigenvalue bound becomes a condition on the flow-topology parameter r of Eq. (15): the state is realisable if the largest strain eigenvalue does not exceed $C_{cr3}|\Omega|/2$, which holds for every strain shape when $r \ge (1 + \sqrt{3}\, C_{cr3}/2)^{-1} = 0.35$ and for none when $r < (1 + \sqrt{3}\, C_{cr3})^{-1} = 0.21$. As shown in Fig. 23d, the WBJ data follow this bound. For all three variants, every state with $r < 0.2$ and 99.5-100% of those with $0.20 \le r < 0.25$ lie outside the triangle, and no variant has a violation beyond $r = 0.35$. The three variants follow nearly the same curve up to $r \approx 0.3$; in the transition band the quadratic closure recovers first (7% of its states outside the triangle for $0.30 \le r < 0.35$, against 38-47% for the linear variants). The violating states lie almost entirely in the outer part of the boundary layer (99% at $y/T > 0.2$), where the mean shear is weak and the strain of the flow around the wing dominates. There the proxy, not the closure, decides the outcome, because without a transported TKE the realisability of an eddy-viscosity closure in strain-dominated flow cannot be assessed.

The triangles of Fig. 23a-c are therefore drawn for the shear- and rotation-dominated states, $r \ge 0.45$, where the proxy is meaningful: the inner part of the plate boundary layer, the corner region and the horseshoe vortex (∼50k states, 57% of the turbulent states). The linear variants are realisable on this population by construction, because $r \ge 0.45$ exceeds the bound of 0.35. Their states spread around the pure-shear point (median $\eta \approx 0.090$) because the gradient is three-dimensional, the rotation-dominated states approaching the isotropic point and ξ taking both signs with $\det(S_{ij})$. The quadratic closure has no such guarantee and is realisable in 100% of these states; it shifts them to the axisymmetric-expansion side, as in the duct (89% at $\xi > 0$, median $(\xi, \eta) = (0.096, 0.123)$, maximum $\eta = 0.156$). The realisability of the calibrated nonlinear closure therefore survives the step from the quasi-two-dimensional duct to a three-dimensional junction, within the region where the $k \propto |\Omega|$ proxy can

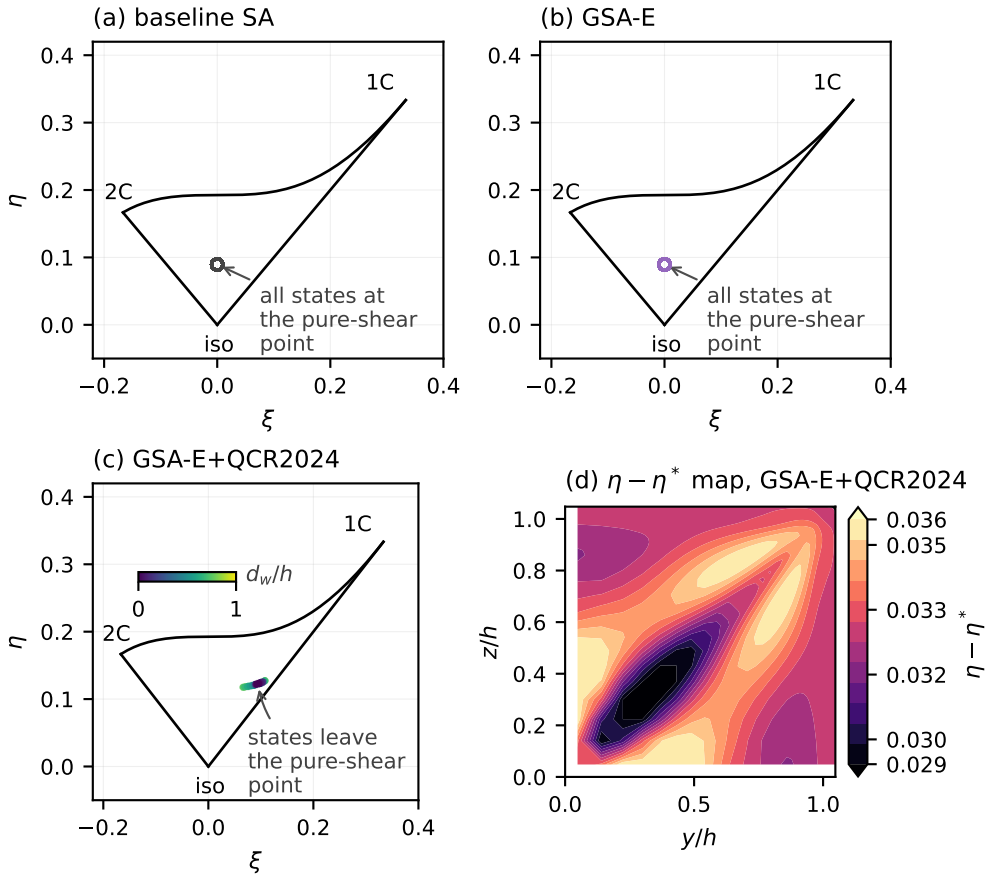

Fig. 22 Anisotropy-invariant map of the modelled Reynolds stress in the square duct, all cells of the cross-section except the wall-adjacent layer, where ν_t and the proxy vanish; TKE proxy $k = \frac{3}{2} C_{cr3} \nu_t |\Omega|$ with the published $C_{cr3} = 2.15$. (a) baseline SA and (b) GSA-E: every cell falls on the pure-shear point $(\xi, \eta) = (0,\ 0.090)$. (c) GSA-E+QCR2024, coloured by wall distance: the states move from the pure-shear point to the axisymmetric-expansion side of the map, around the pure-shear state of the quadratic closure, $(\xi, \eta) \approx (0.094, 0.122)$, and remain well inside the triangle. (d) Contours of the excess anisotropy $\eta - \eta^*$ of GSA-E+QCR2024.

test it. The topology parameter r of the shielding marks this region a priori, which supports the topology-gated closure proposed in the outlook.

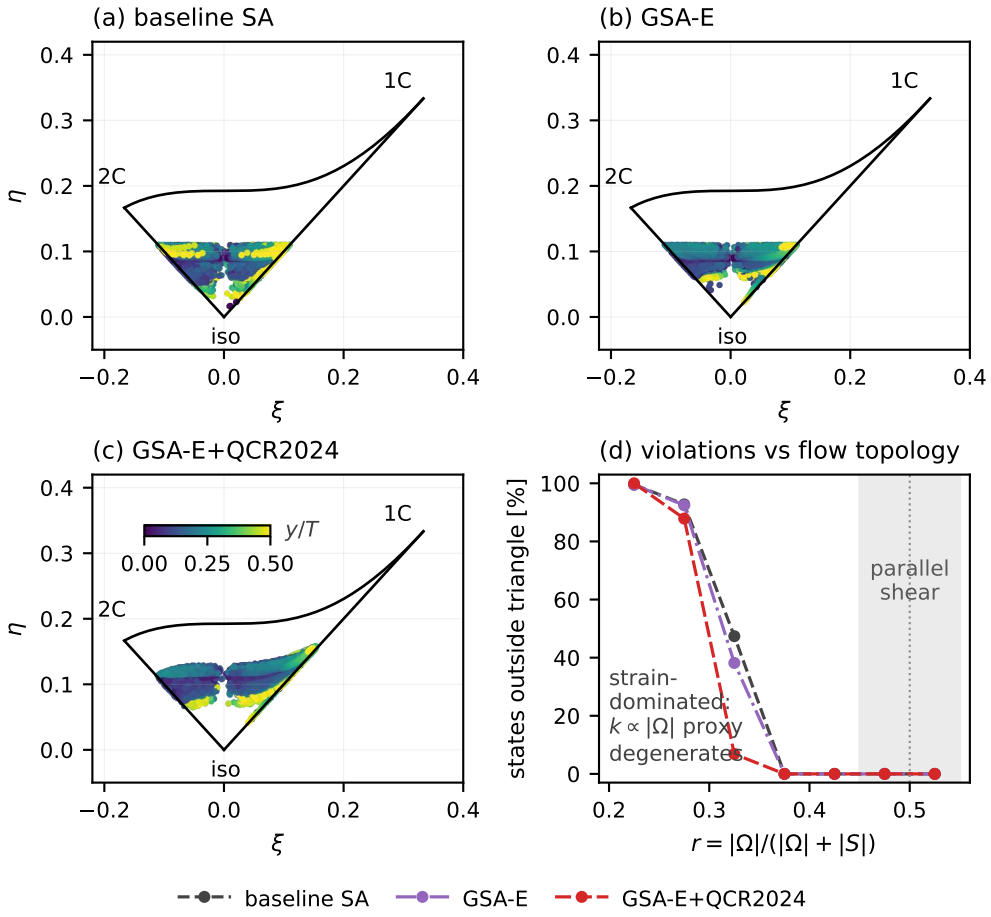

Fig. 23 Anisotropy-invariant map on the WBJ, sampled on a lattice around the junction; TKE proxy as in Fig. 22. (a) baseline SA, (b) GSA-E and (c) GSA-E+QCR2024: shear- and rotation-dominated turbulent states ($r \geq 0.45$, $\nu_t > 5\nu$), coloured by the plate-normal distance y/T; for the linear variants the selection caps η at 0.109. (d) Fraction of states outside the triangle per bin of the flow-topology parameter $r = |\Omega|/(|\Omega| + |S|)$ (the states with $r < 0.2$, all outside, are not shown): violations are confined to strain-dominated states, where the proxy $k \propto |\Omega|$ degenerates.

5 Conclusions and outlook

We have generalised the SA model by giving each calibration coefficient an Inner (near-wall) and an Outer (free-shear) value, blended pointwise by a DDES/IDDES-type shielding function, and calibrated both states by multi-objective optimisation with every objective normalised by the error of baseline SA. The transport equation and its numerics are unchanged; only the calibration space is extended.

Three messages emerge. First, the generalisation does not compromise the SA core. The C_{w1} reconstruction keeps the Inner state log-law consistent, and the attached cases are insensitive to it within the calibration tolerance, in the native suite, in the

other suite and in the attached flow over the unseen aerofoil. Second, the gains, the conflicts and the transferability are all set by the Outer state. A free jet requires the Outer production below the baseline value, a separated shear layer well above it. The NASA calibration contains both and reconciles them: GSA-N reduces the jet objective twentyfold and the hump objective by a fifth. The ERCOFTAC calibration contains no free jet to bound it and lets the Outer production of GSA-E grow to four times that of GSA-N, so GSA-E fails on the NASA jet and hump and suppresses stall on the aerofoil within the computed range, a delay that the quadratic closure removes. GSA-N, instead, improves three of the four ERCOFTAC cases, keeps the stall onset of baseline SA and reduces the junction velocity errors by 15-18%; it is the member of choice when the target flow is not corner-dominated. Third, the calibration separates the deficits that coefficients can fix from those that are structural. No coefficient set generates the corner secondary flow of a square duct, because it is a Reynolds-stress anisotropy effect beyond the reach of any linear eddy-viscosity closure; even a duct-dedicated calibration removes only 20% of the baseline error. A topology-sensitive floor on the separation sensor relaxes the resulting gentle-hill/square-duct trade-off and isolates how much the shielding function alone can reach. A calibrated QCR2024 closure supplies the missing anisotropy: it reduces the square-duct objective of the selected member by 85% with respect to the blending-only model, generates the corner-vortex pair through its rotation term and also lowers both hill costs, at a cost on the step. Its modelled stress stays realisable in every duct cell outside the wall-adjacent layer and in all the shear- and rotation-dominated states of the junction, a non-trivial result for QCR2024, whereas for the linear variants the TKE proxy guarantees realisability in these states by construction (Section 4.10).

Each campaign is a single realisation of the stochastic search. The narrow bounds of run1 keep every SA term physically admissible and already yield members that improve on baseline SA on all four cases of each suite, with smaller gains; the widened bounds go beyond, and the selected members, with $D_I < 0$ and, for GSA-E, $C_{w1I} < 0$ (Sections 3.3 and 4.3), are effective parameters within the tested envelope, not physically interpretable SA constants. On the hump, only the dedicated optimum recovers the measured skin friction downstream of reattachment, and the steep hill stays marginally above unity for GSA-E, because the two hills require opposite Outer productions. The stabilised QCR2024 formulation makes the nonlinear closure usable on external-aerodynamic meshes, but it relies on a pseudo-time relaxation and a stress limiter, which acts only mildly on the aerofoil; its main benefit lies in corner-flow regions, which the junction measurements sample only indirectly, and the published constants make the gentle-hill shear stress oscillate. Finally, the flow-topology sensor uses $|\Omega|$, which is not frame-invariant.

Natural next steps are a calibration on the merged suites, starting from GSA-N, possibly with $J_i \leq 1$ enforced on every case and with an incipient-separation high-lift case bounding the Outer state from the attachment side; a frame-invariant separation sensor; and the promotion of three-dimensional corner cases such as the WBJ from validation to calibration. The most promising structural extension is a topology-gated quadratic closure. The Inner/Outer axis cannot distinguish a duct corner from an attached aerofoil boundary layer (both have $\mathcal{F}_{SA} \to 1$), whereas the flow-topology

sensor r of Eq. (15) can, and it also marks where the vorticity scaling of the closure is trustworthy. Modulating the QCR coefficients, and possibly the Outer production, by a function of r would activate the anisotropic stress where corner physics demands it, let the mixing depend on the type of separation and withdraw both from the stiff attached regions, a physics-based alternative to the stress limiter of Section 2.3. We release the generalised model, the topology floor and the calibration harness as open-source software to support these extensions.

These results bear on a broader question: should a RANS closure be one general model for every flow, or a family of coefficient sets tuned to flow classes? The single-case optima show how much per-flow tuning can reach, and the cross-suite tests how little of it transfers. The present model takes an intermediate position, one formulation whose coefficients change with the local flow regime, so that every gain can be traced to a coefficient and a mechanism. A data-rich, possibly learning-based selection layer could choose the coefficient set suited to each application and remain explainable, because it acts on coefficients with a known role; we nevertheless see the long-term aim of turbulence modelling in physically grounded, interpretable models that leave less to selection.

Nomenclature

symbol	meaning
$\tilde{\nu}$	SA modified eddy viscosity
ν_t, ν_{eff}	eddy viscosity, effective viscosity $\nu + \nu_t$
$\mathcal{F}_{SA}$	shielding (blending) function $\in [0, 1]$
ψ_I, ψ_O	Inner / Outer value of a blended coefficient
$\sigma_{\tilde{\nu}}, C_{b1}, C_{b2}$	blended SA coefficients (diffusion, production, gradient diffusion)
$D = (1 + C_{b2})/\sigma_{\tilde{\nu}}$	front-spreading coefficient of the diffusion term (in x/D, r/D: jet nozzle diameter)
C_{w1}	SA wall-destruction coefficient (rebuilt, Eq. (3))
$r_d, f_d, f_{d\beta}$	DDES wall-distance shielding, wall-distance branch of $\mathcal{F}_{SA}$
f_{step}, Δ	IDDES grid-step blend, IDDES length scale (Section 2.1)
$l_\Omega, x_\Omega, \text{Re}_\Omega, f_{d\Omega}$	vorticity-Reynolds branch quantities, vorticity branch of $\mathcal{F}_{SA}$
r_{vc}, r_{vc0}, C_{rvc}	quiescent floor, baseline floor, topology gain
r	flow-topology parameter $\lvert\Omega\rvert/(\lvert\Omega\rvert + \lvert S\rvert)$
$\lvert S\rvert, \lvert\Omega\rvert, \lvert\nabla U\rvert$	strain, vorticity and velocity-gradient magnitudes
$C_{cr1}, C_{cr2}, C_{cr3}$	QCR2024 constants (rotation, quadratic, isotropic)
O_{ij}, S^*_{ij}, W_{ij}	normalised rotation, deviatoric strain, rotation tensors
$\alpha_{\text{QCR}}, \lambda_{\text{QCR}}$	QCR relaxation factor and stress-limiter ratio
$e_{i,m}, e^{SA}_{i,m}$	error of term m of case i and its baseline-SA value (Eq. (19))
J_i, ϵ	normalised cost of case i, error floor (Eq. (20))
b_{ij}	anisotropy tensor of the modelled Reynolds stress
ξ, η, η^*	anisotropy invariants (Lumley coordinates), pure-shear value of η
u_s, U_b	in-plane (secondary) velocity magnitude, bulk velocity
c_{fd}, c_d, κ, d	shielding constants, von Kármán constant, wall distance

acronym	meaning
SA, GSA	Spalart-Allmaras model, generalised SA model of this work
GSA-N, GSA-E, GSA-E+QCR2024	GSA calibrated on the NASA suite, on the ERCOFTAC suite, with QCR2024
run1-run4	the four multi-case calibration campaigns (Section 3.3)
CFD	computational fluid dynamics
RANS, LES, DNS	Reynolds-averaged, large-eddy, direct numerical simulation

rms	root-mean-square
DDES, IDDES	delayed and improved delayed detached-eddy simulation
QCR	quadratic constitutive relation (QCR2000, QCR2024)
PSO, NSGA-II	particle-swarm optimisation, non-dominated sorting genetic algorithm II
TMR	NASA Turbulence Modeling Resource
ERCOFTAC	European Research Community on Flow, Turbulence and Combustion
2DZP, 2DFDC, ASJ, 2DWMH	TMR flat plate, high-Re channel, axisymmetric subsonic jet, wall-mounted hump
2DN00, ATB, WBJ	NACA 0012 aerofoil, axisymmetric transonic bump, wing-body junction
CBFS	curved backward-facing step
TKE	turbulent kinetic energy

Acknowledgements. The authors thank Leonardo Labs for the computing time on the davinci-1 supercomputer and the HPC@PoliTO academic computing centre of Politecnico di Torino for the computational resources provided.

Declarations

Funding. The authors received no specific funding for this work.

Competing interests. The authors declare no competing interests.

Ethics approval. Not applicable.

Use of AI-assisted tools. During the preparation of this work, the authors used Claude (Anthropic) to assist with post-processing scripts, figure preparation and language editing. The authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Data and code availability. All software and data behind this work are released together. The generalised model is implemented as a single OpenFOAM-13 library, `libGSA`, with incompressible and compressible variants; it includes the topology floor of the shielding function (Eq. (13)), the QCR2024 relation with the optional stabilisation of Section 2.3, and the small HiSA coupling used for the transonic bump. The calibration harness, `gsa_opt`, is a Python package that couples the PSO and NSGA-II optimisers to the OpenFOAM case set-ups of the two suites and logs every error term of every evaluation. The release contains (i) the library and the harness with their documentation; (ii) the case set-ups of the eight calibration cases and of the three unseen configurations, with meshes, solver settings, discretisation schemes and iteration budgets; (iii) the reference data of the calibration cases and the post-processing conventions of every error term; (iv) the single-case optima, the Pareto fronts of run1-run4 with their per-evaluation cost decompositions, and the coefficient sets of GSA-N, GSA-E and GSA-E+QCR2024; (v) the retained field samples of these members and of baseline SA behind every figure; and (vi) the `numpy`/`matplotlib` scripts that regenerate all figures and tables from the archived files, with no proprietary dependencies, so that every number in Section 4 can be reproduced. The package will be deposited on Zenodo upon acceptance of the manuscript under the reserved persistent identifier doi:10.5281/zenodo.22836636; the code is released under the GPL-3.0 licence and the

data under CC-BY-4.0. Until the deposit, the package is available to the editor and the referees on request.

Author contributions. T. Di Fabbio: conceptualisation, methodology, software (model library and calibration harness), investigation, formal analysis, visualisation, writing - original draft. R. Vallino: software, validation, investigation. E. Segalerba: methodology, validation, writing - review and editing. V. Rossano: resources, investigation, writing - review and editing. J. Guerrero: resources, supervision, writing - review and editing. A. Ferrero: conceptualisation, supervision, project administration, writing - review and editing. All authors read and approved the final manuscript.

References

Allmaras SR, Johnson FT, Spalart PR (2012) Modifications and clarifications for the implementation of the Spalart–Allmaras turbulence model. In: 7th International Conference on Computational Fluid Dynamics (ICCFD7), Big Island, Hawaii, paper ICCFD7-1902

Bachalo WD, Johnson DA (1986) Transonic, turbulent boundary-layer separation generated on an axisymmetric flow model. AIAA Journal 24(3):437–443. https://doi.org/10.2514/3.9286

Bentaleb Y, Lardeau S, Leschziner MA (2012) Large-eddy simulation of turbulent boundary layer separation from a rounded step. Journal of Turbulence 13(4):1–28. https://doi.org/10.1080/14685248.2011.637923

Breuer M, Peller N, Rapp C, et al (2009) Flow over periodic hills - numerical and experimental study in a wide range of Reynolds numbers. Computers & Fluids 38(2):433–457. https://doi.org/10.1016/j.compfluid.2008.05.002

Bridges J, Wernet MP (2010) Establishing consensus turbulence statistics for hot subsonic jets. In: 16th AIAA/CEAS Aeroacoustics Conference, Stockholm, Sweden, https://doi.org/10.2514/6.2010-3751, AIAA Paper 2010-3751

Buchanan T, Cinnella P, Dwight RP, et al (2026) Data-driven turbulence modelling challenge, 3rd ERCOFTAC workshop on machine learning for fluid dynamics. https://ml4fluids2026.github.io/challenges/, Amsterdam, 4–6 March 2026; accessed September 2026

Cherroud S, Merle X, Cinnella P, et al (2022) Sparse Bayesian learning of explicit algebraic Reynolds-stress models for turbulent separated flows. International Journal of Heat and Fluid Flow 98:109047. https://doi.org/10.1016/j.ijheatfluidflow.2022.109047

Deb K, Pratap A, Agarwal S, et al (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation 6(2):182–197. https://doi.org/10.1109/4235.996017

Demuren AO, Rodi W (1984) Calculation of turbulence-driven secondary motion in non-circular ducts. Journal of Fluid Mechanics 140:189–222. https://doi.org/10.1017/S0022112084000574

Devenport WJ, Simpson RL (1990) Time-dependent and time-averaged turbulence structure near the nose of a wing-body junction. Journal of Fluid Mechanics 210:23–55. https://doi.org/10.1017/S0022112090001215

Di Fabbio T, Fang Y, Tangermann E, et al (2024) Strategies for enhancing one-equation turbulence model predictions using gene-expression programming. Fluids 9(8):191. https://doi.org/10.3390/fluids9080191

Duraisamy K, Iaccarino G, Xiao H (2019) Turbulence modeling in the age of data. Annual Review of Fluid Mechanics 51(1):357–377. https://doi.org/10.1146/annurev-fluid-010518-040547

Edeling WN, Cinnella P, Dwight RP (2014a) Predictive RANS simulations via Bayesian model-scenario averaging. Journal of Computational Physics 275:65–91. https://doi.org/10.1016/j.jcp.2014.06.052

Edeling WN, Cinnella P, Dwight RP, et al (2014b) Bayesian estimates of parameter variability in the k–ε turbulence model. Journal of Computational Physics 258:73–94. https://doi.org/10.1016/j.jcp.2013.10.027

Fang Y, Zhao Y, Waschkowski F, et al (2023) Toward more general turbulence models via multicase computational-fluid-dynamics-driven training. AIAA Journal 61(5):2100–2115. https://doi.org/10.2514/1.J062572

Graftieaux L, Michard M, Grosjean N (2001) Combining PIV, POD and vortex identification algorithms for the study of unsteady turbulent swirling flows. Measurement Science and Technology 12(9):1422–1429. https://doi.org/10.1088/0957-0233/12/9/307

Greenblatt D, Paschal KB, Yao CS, et al (2006) Experimental investigation of separation control part 1: Baseline and steady suction. AIAA Journal 44(12):2820–2830. https://doi.org/10.2514/1.13817

Heyns JA, Oxtoby OF, Steenkamp A (2014) Modelling high-speed flow using a matrix-free coupled solver. In: 9th OpenFOAM Workshop, 23–26 June 2014, Zagreb, Croatia, HiSA (High Speed Aerodynamic solver), https://gitlab.com/hisa/hisa

Kennedy J, Eberhart R (1995) Particle swarm optimization. In: Proceedings of ICNN 95 - International Conference on Neural Networks, pp 1942–1948, https://doi.org/10.1109/ICNN.1995.488968

Kotzlowski D, Di Fabbio T, Tangermann E, et al (2026) Enhancing one-equation turbulence models for delta wings by gene expression programming. AIAA Journal

64(4):2189–2205. https://doi.org/10.2514/1.J065906

Ladson CL (1988) Effects of independent variation of Mach and Reynolds numbers on the low-speed aerodynamic characteristics of the NACA 0012 airfoil section. Tech. Rep. NASA TM-4074, NASA Langley Research Center

Ling J, Kurzawski A, Templeton J (2016) Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. Journal of Fluid Mechanics 807:155–166. https://doi.org/10.1017/jfm.2016.615

Liu C, Wang Y, Yang Y, et al (2016) New omega vortex identification method. Science China Physics, Mechanics & Astronomy 59(8):684711. https://doi.org/10.1007/s11433-016-0022-6

Lumley JL (1979) Computational modeling of turbulent flows. Advances in Applied Mechanics 18:123–176. https://doi.org/10.1016/S0065-2156(08)70266-7

Margheri L, Meldi M, Salvetti MV, et al (2014) Epistemic uncertainties in RANS model free coefficients. Computers & Fluids 102:315–335. https://doi.org/10.1016/j.compfluid.2014.06.029

McConkey R, Yee E, Lien FS (2021) A curated dataset for data-driven turbulence modelling. Scientific Data 8(1):255. https://doi.org/10.1038/s41597-021-01034-2

McConkey R, Buchanan T, Smidt T, et al (2026) The closure challenge: A benchmark task for machine learning in turbulence modelling. https://doi.org/10.48550/arXiv.2603.28884, arXiv preprint 2603.28884

McCroskey WJ (1987) A critical assessment of wind tunnel results for the NACA 0012 airfoil. Tech. Rep. NASA TM-100019, NASA Ames Research Center

Menter FR (1994) Two-equation eddy-viscosity turbulence models for engineering applications. AIAA Journal 32(8):1598–1605. https://doi.org/10.2514/3.12149

Menter FR, Matyushenko A (2025) Generalized k–ω (GEKO) two-equation turbulence model. AIAA Journal 63(11):4590–4606. https://doi.org/10.2514/1.J065393

Parish EJ, Duraisamy K (2016) A paradigm for data-driven predictive modeling using field inversion and machine learning. Journal of Computational Physics 305:758–774. https://doi.org/10.1016/j.jcp.2015.11.012

Pope SB (1978) An explanation of the turbulent round-jet/plane-jet anomaly. AIAA Journal 16(3):279–281. https://doi.org/10.2514/3.7521

Rumsey CL (2026) NASA turbulence modeling resource (TMR). https://tmbwg.github.io/turbmodels/, maintained by the NASA Langley Research Center and the Turbulence Model Benchmarking Working Group; accessed September 2026

Rumsey CL, Smith BR, Huang GP (2010) Description of a website resource for turbulence modeling verification and validation. In: 40th AIAA Fluid Dynamics Conference and Exhibit, https://doi.org/10.2514/6.2010-4742, AIAA Paper 2010-4742

Rumsey CL, Coleman GN, Wang L (2022) In search of data-driven improvements to RANS models applied to separated flows. In: AIAA SciTech 2022 Forum, https://doi.org/10.2514/6.2022-0937, AIAA Paper 2022-0937

Rung T, Bunge U, Schatz M, et al (2003) Restatement of the Spalart-Allmaras eddy-viscosity model in strain-adaptive formulation. AIAA Journal 41(7):1396–1399. https://doi.org/10.2514/2.2089

Schaefer J, Hosder S, West T, et al (2017) Uncertainty quantification of turbulence model closure coefficients for transonic wall-bounded flows. AIAA Journal 55(1):195–213. https://doi.org/10.2514/1.J054902

Schmelzer M, Dwight RP, Cinnella P (2020) Discovery of algebraic Reynolds-stress models using sparse symbolic regression. Flow, Turbulence and Combustion 104(2–3):579–603. https://doi.org/10.1007/s10494-019-00089-x

Schumann U (1977) Realizability of Reynolds-stress turbulence models. Physics of Fluids 20(5):721–725. https://doi.org/10.1063/1.861942

Shur ML, Strelets MK, Travin AK, et al (2000) Turbulence modeling in rotating and curved channels: Assessing the Spalart–Shur correction. AIAA Journal 38(5):784–792. https://doi.org/10.2514/2.1058

Shur ML, Spalart PR, Strelets MK, et al (2008) A hybrid RANS–LES approach with delayed-DES and wall-modelled LES capabilities. International Journal of Heat and Fluid Flow 29(6):1638–1649. https://doi.org/10.1016/j.ijheatfluidflow.2008.07.001

Singh AP, Medida S, Duraisamy K (2017) Machine-learning-augmented predictive modeling of turbulent separated flows over airfoils. AIAA Journal 55(7):2215–2227. https://doi.org/10.2514/1.J055595

Spalart PR (2000) Strategies for turbulence modelling and simulations. International Journal of Heat and Fluid Flow 21(3):252–263. https://doi.org/10.1016/S0142-727X(00)00007-2

Spalart PR, Allmaras SR (1992) A one-equation turbulence model for aerodynamic flows. In: 30th Aerospace Sciences Meeting and Exhibit, https://doi.org/10.2514/6.1992-439, AIAA Paper 92-0439

Spalart PR, Deck S, Shur ML, et al (2006) A new version of detached-eddy simulation, resistant to ambiguous grid densities. Theoretical and Computational Fluid Dynamics 20(3):181–195. https://doi.org/10.1007/s00162-006-0015-0

Tamaki Y, Kawai S (2024) Turbulence anisotropy effects on corner-flow separation: Physics and turbulence modelling. Journal of Fluid Mechanics 980:A21. https://doi.org/10.1017/jfm.2024.25

Vinuesa R, Noorani A, Lozano-Durán A, et al (2014) Aspect ratio effects in turbulent duct flows studied through direct numerical simulation. Journal of Turbulence 15(10):677–706. https://doi.org/10.1080/14685248.2014.925623

Waschkowski F, Zhao Y, Sandberg R, et al (2022) Multi-objective CFD-driven development of coupled turbulence closure models. Journal of Computational Physics 452:110922. https://doi.org/10.1016/j.jcp.2021.110922

Weatheritt J, Sandberg RD (2016) A novel evolutionary algorithm applied to algebraic modifications of the RANS stress–strain relationship. Journal of Computational Physics 325:22–37. https://doi.org/10.1016/j.jcp.2016.08.015

Weller HG, Tabor G, Jasak H, et al (1998) A tensorial approach to computational continuum mechanics using object-oriented techniques. Computers in Physics 12(6):620–631. https://doi.org/10.1063/1.168744

Xiao H, Cinnella P (2019) Quantification of model uncertainty in RANS simulations: A review. Progress in Aerospace Sciences 108:1–31. https://doi.org/10.1016/j.paerosci.2018.10.001

Xiao H, Wu JL, Laizet S, et al (2020) Flows over periodic hills of parameterized geometries: A dataset for data-driven turbulence modeling from direct simulations. Computers & Fluids 200:104431. https://doi.org/10.1016/j.compfluid.2020.104431

Zhao Y, Akolekar HD, Weatheritt J, et al (2020) RANS turbulence model development using CFD-driven machine learning. Journal of Computational Physics 411:109413. https://doi.org/10.1016/j.jcp.2020.109413