



Elliot James Badcock · Shahid Mughal

A carrier-wave factored one-way Navier–Stokes method for boundary-layer instability modelling

Received: 16 November 2025 / Accepted: 23 June 2026
© The Author(s) 2026

Abstract We present M-OWNS, a spatial marching method that combines the carrier-wave factoring of the parabolised stability equations (PSE) with a recursive one-way Navier–Stokes (OWNS-R) projection framework. A distinct numerical resolution and efficiency advantage is offered by the approach, in modelling disturbance and instability state evolution. A spectral resolution comparison analysis shows that to leading order, for any excited eigenfunction whose eigenvalue lies closer to the carrier wavenumber than to the origin, the wave-factored system resolves the mode at a coarser streamwise numerical step size relative to the unfactored system. A non-iterating variant, with the carrier wavenumber determined from the base flow, temporal frequency and spanwise wavenumber alone, achieves equivalent resolution accuracy at identical per-step cost to unfactored OWNS. For the fixed-carrier variant, M-OWNS reduces the total solve count by factors of two to eight relative to unfactored OWNS across the test cases considered, with larger reductions possible when the iterated closure condition of PSE is suitable. The method is validated across incompressible and subsonic flat-plate boundary-layers, three-dimensional crossflow disturbances, and a Mach 4.5 hypersonic boundary-layer with four forcing configurations: eigenfunction inlet forcing, wall suction/blowing, multi-mode freestream forcing and randomised inlet forcing. The wall suction/blowing case is validated against a fully elliptic linear harmonic Navier–Stokes solver. For deterministic forcing scenarios, M-OWNS captures disturbance amplitudes, acoustic radiation fields, and modal synchronisation sequences at coarser streamwise resolution than unfactored OWNS. Under broadband randomised forcing, M-OWNS resolves mixed-mode disturbance development at half the numerical cost relative to standard OWNS.

Keywords Boundary-layer instability · Parabolised stability equations (PSE) · One-Way Navier–Stokes (OWNS) · Carrier-wave factoring

1 Introduction

Predicting laminar–turbulent transition in boundary layers requires models that can track disturbance evolution from receptivity through amplification across a developing base flow. The simplest method, local eigenvalue analysis (LST), assumes parallel flow and examines one streamwise station at a time [1]. At the other extreme, global methods solve the full elliptic linearised harmonic Navier–Stokes equations (LHNS), capturing both upstream and downstream disturbance dynamics without invoking a parallel-flow assumption [2,3]. The generality comes at substantial cost, principally through memory constraints and the need to invert large matrices. For convective disturbances, where upstream and downstream dynamics evolve independently, spatial marching

E. J. Badcock (✉) · S. Mughal
Department of Mathematics, Imperial College London, South Kensington Campus, SW7 2AZ London, United Kingdom
E-mail: ejb321@ic.ac.uk

S. Mughal
E-mail: s.mughal@imperial.ac.uk

methods offer an attractive alternative: the solution is advanced station by station in the downstream direction at a fraction of the global cost. The difficulty is that the LHNS operator supports upstream-propagating acoustic modes alongside the downstream-propagating disturbances of interest. These upstream modes can destabilise a naïve spatial march; practical methods must either remove, suppress or neglect them at coarser resolutions.

The parabolised stability equations (PSE) contend with upstream modes through a WKB-like wave-carrier factoring: the disturbance is decomposed into a rapidly varying carrier wave and a slowly varying envelope, with the carrier wavenumber α_0 updated iteratively at each streamwise station [4,5]. The factoring shifts the downstream excited spectrum toward the origin, so the physically relevant modes can be resolved at coarse streamwise resolution. For single mode disturbances this is highly effective; the iteration drives α_0 toward the wavenumber of the dominant mode, stretching its wavelength. However, the upstream modes remain in the system. At coarse resolution they fall outside the integrator's accuracy region and cause no harm, but when fine resolution is needed (multi-modal/non-modal disturbances, receptivity near the leading edge [6], nonlinear calculations) the accuracy region widens and captures them, destabilising the numerical marching procedure [7,8]. Regularisation strategies (implicit damping *e.g.* backward Euler marching [9] or neglecting the streamwise pressure gradient [4]) permit finer steps but introduce characteristic errors in the wave ansatz and shape function. Towne et al. [10] trace these errors to the step-size restriction and the regularisation required to enforce it. To our knowledge, the wave-carrier factoring itself has never been tested independently of the restriction within a spatial marching method.

The One-Way Navier–Stokes (OWNS) method takes a different approach: rather than shifting eigenvalues away from the upstream modes, it removes the upstream modes entirely via one-way projection [11,12], thus eliminating the PSE step-size restriction, and importantly retaining the streamwise pressure gradient. The projection is controlled by recursion parameters placed in the complex plane, which determine which spectral regions the projection removes or retains. These parameters are specific to the flow configuration and forcing frequency, and must be selected with care. The OWNS-R variant [13] made the projection computationally tractable by evaluating it through a sequence of recursions. Despite its success in parabolisation, OWNS does not employ wave-carrier factoring. The discretisation must resolve eigenvalues at their full wavenumbers from the origin, and accurate prediction of disturbance amplitudes requires fine streamwise resolution. Fine resolution in turn widens the integrator's spectral window, demanding that the projection suppresses a broader portion of the upstream-propagating spectrum. The OWNS-R recursion has a practical ceiling on the number of recursion parameters N_β , beyond which numerical error accumulates [14], so a wider spectral window requires more carefully placed parameters, a non-trivial task for general flows.

The present paper combines the two approaches. The carrier-wave factoring of PSE is applied within the OWNS-R framework, with the carrier wavenumber α_0 chosen from the eigenvalue spectrum of the spatial operator or from a representative value of the excited modes; this spectral anchoring is what makes the factoring computationally beneficial (§2.4). We retain the name M-OWNS from Badcock and Mughal [15]: the M refers to this choice of wave-carrier being tied to the excited spectrum of the operator, not to any modal restriction on the disturbance. Unfactored OWNS corresponds to $\alpha_0 = 0$. The one-way projection removes the upstream-propagating modes that would otherwise destabilise a spatial march in PSE, so the factoring benefit is realised without the minimum step-size restriction of PSE.

A leading-order spectral resolution comparison analysis (§2.4) shows that, for any excited eigenfunction whose eigenvalue lies closer to α_0 than to the origin, the wave-factored system resolves that mode at a coarser step size than the unfactored system. The spectral geometry of boundary-layer flows supports this in two complementary regimes. When the dominant response is composed of discrete modes (Tollmien–Schlichting (T–S), crossflow, or slow and fast hypersonic waves) clustered near the vorticity/entropy continuous-spectrum branch endpoint α_1 of Eq. (B.1), a carrier at or near α_1 recentres the numerical scheme on the excited cluster. When several branches are simultaneously excited, such as in receptivity or optimal-disturbance calculations, a carrier placed at a representative centre of the excited spectrum covers the full range at once (§2.5). For single-mode calculations, α_0 can be iterated toward the dominant eigenvalue as in PSE. For multi-modal or receptivity calculations, fixing α_0 at a suitable spectral centre avoids iteration entirely at identical per-step cost to unfactored OWNS.

The carrier-wave factoring approach is presented within the OWNS-R framework, but the improved spectral resolution achieved is not specific to OWNS. We use OWNS-R because the one-way projection permits fine spatial marching steps without the regularisation errors PSE incurs [10]. For the fixed-carrier variant with the second-order implicit (BFD2) scheme, M-OWNS reduces the total solve count by factors of two to eight relative to unfactored OWNS across the test cases considered; larger reductions follow when a single mode dominates and the Rayleigh-quotient iteration applies (§2.5).

Section 2 introduces the governing equations and wave-carrier formulation, and §3 describes the OWNS-R projection and M-OWNS algorithm. The method is validated on flows of increasing complexity: subsonic flat-plate boundary layers (§5.1), three-dimensional crossflow over a swept cylinder [16] (§5.2), and a Mach 4.5 hypersonic boundary layer [1] (§5.3). The hypersonic case is tested under forcing conditions ranging from single eigenfunction inlet forcing to broadband randomised forcing, including fast acoustic freestream waves and wall suction/blowing. For each deterministic forcing configuration, M-OWNS at coarse streamwise resolution captures disturbance amplitudes, acoustic radiation fields, and modal synchronisation sequences that unfactored OWNS at the same resolution does not. The randomised forcing case demonstrates convergence under broadband excitation, with M-OWNS and OWNS initially resolving complementary parts of the spectrum as resolution is increased. In both iterated and non-iterated forms, we demonstrate that M-OWNS captures the full field at a coarser spatial resolution and reduced computational cost relative to conventional OWNS.

2 Governing equations and wave-carrier formulation

We consider compressible boundary-layer flows governed by the body-fitted Navier–Stokes equations summarised in Appendix A. The base flow is assumed invariant in the spanwise (z) direction, as appropriate for infinite swept wings. Technical details of the OWNS-R approach are given by Zhu and Towne [13]; the present section introduces only the governing equations and the carrier-wave factoring that underpins M-OWNS.

2.1 Two-way disturbance evolution equation

We begin by decomposing the state vector into Q , a known invariant in z and time-independent base flow, and q , a small harmonic disturbance with frequency ω and a spanwise Fourier mode of wavenumber β in z ,

$$q^{\text{tot}}(x, y, z, t) = Q(x, y) + q(x, y) \exp(i\beta z - i\omega t). \quad (1)$$

Substituting this ansatz into the Navier–Stokes equations (Eqs. (A.3)–(A.9)) and then linearising with respect to q yields the LHNS, which can be expressed as

$$\mathbf{A} \left(\frac{\partial}{\partial y}; \beta \right) \frac{\partial q}{\partial x} = \mathbf{B} \left(\frac{\partial}{\partial y}; \beta, \omega \right) q + \mathbf{C} \frac{\partial^2 q}{\partial x^2}. \quad (2)$$

Here, $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ represent linear operators of size 5×5 . The steady base flow $Q(x, y)$ and its derivatives are obtained from a standard boundary-layer solver [17].

The second-order $\mathbf{C}$ terms arise from the streamwise viscous derivatives in the Navier–Stokes equations. For linear disturbances, these terms may be omitted without loss of accuracy [18]; Zhu and Towne [13] confirm this through extensive numerical testing across a variety of flows. In the present work, the $\mathbf{C}$ terms are retained in the implicit marching discretisation but are omitted from the eigenvalue analysis that follows since they do not enter into the definition of the projection operator used in §3.

After discretising in the wall-normal direction with n_y points, the spatial evolution equation takes the semi-discrete form

$$\mathbf{A} \frac{d\phi}{dx} = \mathbf{B}\phi, \quad (3)$$

where $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ are the banded wall-normal discretisations of the continuous operators $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ in Eq. (2), each of order $5n_y$. The terms associated with $\mathbf{C}$ are incorporated into the discretisation stencil and therefore enter the marching scheme implicitly, rather than through augmentation of the state vector.

Since the pencil $(\mathbf{B}, \mathbf{A})$ is regular, the finite and infinite eigenvalues define disjoint spectral sets. The block diagonalisation of Stewart and Sun [19, Chapter VI, Theorem 2.12] then yields a decomposition into deflating subspaces associated with each set. On the finite-eigenvalue subspace, $\mathbf{A}$ is nonsingular, since this subspace excludes the infinite spectrum removed by the deflation. We assume that the eigenvalues of the pencil on this subspace are semisimple, so that the pencil admits a complete set of eigenfunctions. This assumption is used only for the modal-expansion proof in §3.1 that the projection operator parabolises the wave-carrier factored operator; the method itself does not rely on it.

The pencil reduces to the standard eigenvalue problem

$$\mathbf{M}\phi_k = i\lambda_k \phi_k, \quad \mathbf{M} = \mathbf{A}^{-1}\mathbf{B}, \quad (4)$$

where the inverse is taken on the finite spectral subspace, and (ϕ_k, λ_k) are an eigenfunction and finite eigenvalue pair. Let the columns of $\mathbf{V}$ be the corresponding right eigenfunctions ϕ_k , and let the rows of $\mathbf{U}$ be the associated left eigenfunctions, chosen such that

$$\mathbf{U}\mathbf{V} = \mathbb{I}, \quad (5)$$

on the finite eigenvalue subspace. The spatial dynamics on this subspace, and eigen-decomposition of $\mathbf{M}$, can then be written as

$$\frac{d\phi}{dx} = \mathbf{M}\phi, \quad \mathbf{M} = \mathbf{V} i \mathbf{\Delta} \mathbf{U}, \quad (6)$$

where $\mathbf{\Delta} = \text{diag}(\lambda_k)$, and we write $\sigma(\mathbf{M}) = \{\lambda_k\}$ for the finite eigenvalue spectrum. In practice, $\mathbf{M}$ is not formed explicitly; the numerical methods of §3 operate directly on the pencil $(\mathbf{B}, \mathbf{A})$.

2.2 Disturbance energy and rayleigh quotient

On the semi-infinite wall-normal domain, we define the inner product between two disturbance states ϕ_1 and ϕ_2 as

$$(\phi_1, \phi_2)_{\mathcal{H}} = \int_0^\infty \phi_1^* \mathcal{H} \phi_2 dy, \quad (7)$$

where $(\cdot)^*$ denotes conjugate transpose and $\mathcal{H}$ is a positive semi-definite Hermitian weight matrix. The formulation that follows is independent of the specific choice of $\mathcal{H}$; any positive semi-definite weight matrix may be used. In the present work we adopt the density-weighted kinetic energy norm [20],

$$\mathcal{H} = \text{diag}(0, \rho, \rho, \rho, 0), \quad E_{\mathcal{H}} = \frac{1}{2} \int_0^\infty \rho(|u|^2 + |v|^2 + |w|^2) dy, \quad (8)$$

where $\rho = \rho(x, y)$ is the local base flow density. Results have also been verified with the diagonal part of the compressible energy norm of Chu [21] and Hanifi et al. [22].

Given a solution $\phi(x, y)$ of Eq. (3), we define the Rayleigh quotient

$$\alpha(x) = -i \frac{\left(\phi, \frac{d\phi}{dx} \right)_{\mathcal{H}}}{\|\phi\|_{\mathcal{H}}^2}, \quad (9)$$

where $\|\phi\|_{\mathcal{H}}^2 = (\phi, \phi)_{\mathcal{H}}$.

Interpretation as a wavenumber

The Rayleigh quotient α defined by Eq. (9) does not require solving the eigenvalue problem (4); it is evaluated directly from the disturbance field ϕ and the weight $\mathcal{H}$. Its interpretation as a wavenumber depends on the modal composition of the disturbance.

For a disturbance dominated by a single eigenfunction, $\phi \approx c_k(x) \phi_k(x)$, where $c_k(x)$ is a slowly varying coefficient and (ϕ_k, λ_k) is an eigenfunction/eigenvalue pair of $\mathbf{M}$ (Eq. (4)), the Rayleigh quotient recovers $\alpha = \lambda_k$ exactly when c_k is constant, and approximately when c_k varies slowly relative to the wavelength $2\pi/|\lambda_k|$. The Rayleigh quotient then provides a local streamwise wavenumber and amplification rate.

When several eigenfunctions carry comparable energy, α is a weighted average of the excited eigenvalues. If the eigenfunctions are mutually orthogonal under $\mathcal{H}$, the cross terms $(\phi_j, \phi_k)_{\mathcal{H}}$ vanish and the weights reduce to each mode's share of $\|\phi\|_{\mathcal{H}}^2$. The Rayleigh quotient is then a weighted average of the excited eigenvalues and lies within their convex hull.

When the eigenfunctions are not orthogonal under $\mathcal{H}$, the cross terms do not vanish. They oscillate at the beating wavenumbers $|\lambda_j - \lambda_k|$ and, when multiple non-orthogonal eigenfunctions carry significant amplitude, can dominate over the diagonal contributions. In this regime α may be driven outside the convex hull of the excited eigenvalues and can no longer be interpreted as a weighted average of eigenvalues or as a wavenumber of any individual mode. The Rayleigh quotient should therefore be regarded as a computational reference quantity characterising the dominant phase behaviour of the disturbance, rather than a physically precise wavenumber.

2.3 Carrier-wave factoring

We factor the solution of Eq. (3) as

$$\phi(x) = \tilde{\phi}(x) \exp i\Theta_0(x), \quad \frac{d\Theta_0}{dx} = \alpha_0, \quad (10)$$

where α_0 is a carrier wavenumber and $\tilde{\phi}$ is the envelope. The same decomposition is employed by PSE [4] and adaptive LHNS [23,24]: we refer to it as carrier-wave factoring. As a change of variables it is always formally valid, the only requirement being that α_0 is sufficiently differentiable for the system at hand. The carrier wavenumber α_0 need not equal any eigenvalue in $\sigma(\mathbf{M})$ nor the Rayleigh quotient α defined in Eq. (9); specific choices of α_0 are discussed in §2.5.

Substituting Eq. (10) into the semi-discrete system Eq. (3) yields the factored system

$$\mathbf{A} \frac{d\tilde{\phi}}{dx} = \tilde{\mathbf{B}} \tilde{\phi}, \quad (11)$$

where $\tilde{\mathbf{B}} = \mathbf{B} - i\alpha_0 \mathbf{A}$. The pressure derivative $dp/dx = i\alpha_0 \tilde{p} + d\tilde{p}/dx$ is retained in its entirety; often in PSE the $d\tilde{p}/dx$ term is set to zero in subsonic regions [4].

The factoring shifts the spectrum by a scalar and leaves the eigenfunctions unchanged [10]. The spatial operator of the factored system is

$$\tilde{\mathbf{M}} = \mathbf{V} i(\Delta - \alpha_0 \mathbb{I}) \mathbf{U} = \mathbf{M} - i\alpha_0 \mathbb{I}, \quad (12)$$

where $\mathbf{V}\mathbf{U} = \mathbb{I}$ on the finite-eigenvalue subspace, with shifted spectrum

$$\sigma(\tilde{\mathbf{M}}) = \{\lambda_k - \alpha_0 : \lambda_k \in \sigma(\mathbf{M})\}. \quad (13)$$

When the second-order $\mathbf{C}$ terms are retained, the factored system (11) becomes

$$\tilde{\mathbf{A}} \frac{d\tilde{\phi}}{dx} = \tilde{\mathbf{B}} \tilde{\phi} + \mathbf{C} \frac{d^2 \tilde{\phi}}{dx^2}, \quad (14)$$

where

$$\tilde{\mathbf{A}} = \mathbf{A} - 2i\alpha_0 \mathbf{C}, \quad \tilde{\mathbf{B}} = \mathbf{B} - i\alpha_0 \mathbf{A} + \left(i \frac{d\alpha_0}{dx} - \alpha_0^2 \right) \mathbf{C}. \quad (15)$$

2.4 Spectral resolution comparison

To examine the resolution properties of the wave-factored system, consider the parallel-flow eigenfunction expansion of Eq. (6),

$$\phi(x) = \sum_k v_k \exp(i\lambda_k x), \quad (16)$$

where the eigenfunctions v_k are independent of x . A streamwise integration scheme resolves mode k only if the spatial step is small relative to its wavelength. To leading order this requirement is

$$\Delta x \lesssim \frac{C}{|\lambda_k|}, \quad (17)$$

where C is an order-one constant set by the accuracy region of the scheme, for example $C = \pi$ at the Nyquist limit of a second-order centred scheme. We do not assume a particular scheme. The precise accuracy depends on the position in the complex plane of the *argument* $i\lambda_k \Delta x$, the value at which the scheme is evaluated, and therefore on its phase as well as its modulus, and on the discretisation.

Introducing the carrier wave gives the envelope

$$\tilde{\phi}(x) = \sum_k v_k \exp(i(\lambda_k - \alpha_0)x), \quad (18)$$

whose modes oscillate at the residual wavenumbers $|\lambda_k - \alpha_0|$. The envelope is resolved when

$$\Delta x \lesssim \frac{C}{|\lambda_k - \alpha_0|}. \quad (19)$$

The same constant C appears in Eqs. (17) and (19) because the two systems are advanced with the same discretisation and scheme. The argument of the factored system is $i(\lambda_k - \alpha_0)\Delta x$.

Equation (13) shifts the eigenvalues by $-\alpha_0$; equivalently, holding the eigenvalues fixed, the scheme's resolved region moves to α_0 .

Leading-order

spectral resolution condition Comparing Eqs. (17) and (19) at fixed Δx , factoring relaxes the resolution requirement of mode k whenever

$$|\lambda_k - \alpha_0| < |\lambda_k|. \quad (20)$$

The locus $|\lambda - \alpha_0| = |\lambda|$ is the perpendicular bisector of the segment from 0 to α_0 . Modes on the α_0 side benefit, modes on the far side require a finer step than the unfactored system, and for $|\lambda_k| \gg |\alpha_0|$ the factoring is neutral. The constant C is common to both systems, so the leading-order condition carries no explicit scheme-dependent term.

We call Eq. (20) the spectral resolution condition. It is a leading-order statement. For modes whose argument lies near the centre of the scheme's accuracy region the modulus of the wavenumber alone sets the requirement; nearer the edge the requirement acquires a dependence on the phase of the argument and on the scheme, as noted at Eq. (17), so the true break-even locus is a scheme-dependent curve rather than the bisector. This dependence is common to both factored and unfactored systems and is a property of the integration scheme, not of the factoring.

Spectral clustering

In subsonic and supersonic boundary-layer flows, the downstream-propagating discrete modes and the endpoints of the continuous spectral branches cluster in the complex λ -plane. A carrier placed near this cluster sits close to the excited modes, while the modes far along the branches have large $|\lambda_k|$ and require a fine step with or without factoring. The benefit thus depends on the choice of α_0 . A poorly placed carrier wastes the translation, and since $\alpha_0 = 0$ recovers the unfactored system, the unfactored march is always available as a fallback. Any estimate near the cluster serves as a reasonable α_0 . Specific choices are discussed in §2.5.

Reliability of the comparison

With the carrier near the physically relevant cluster, the modes (λ_k) closest to α_0 satisfy $|\lambda_k - \alpha_0| \ll |\lambda_k|$. For a fixed step size Δx , the factored argument $|(\lambda_k - \alpha_0)\Delta x|$ is then small and lies near the centre of the scheme's accuracy region. In this limit the modulus governs and the phase- and scheme-dependent corrections are subdominant, so the comparison is reliable for these modes.

Modes within the cluster but at finite distance from α_0 need not satisfy the stronger inequality $|\lambda_k - \alpha_0| \ll |\lambda_k|$. They may still satisfy the spectral resolution condition $|\lambda_k - \alpha_0| < |\lambda_k|$, but the argument $|(\lambda_k - \alpha_0)\Delta x|$ is then not necessarily small. The phase- and scheme-dependent corrections therefore enter both the factored and unfactored systems, and unless a specific scheme, spectrum, and choice of α_0 are prescribed, neither system can be claimed better resolved at this order.

As Δx decreases, the arguments of both systems move towards the centre of the scheme's accuracy region. For modes satisfying $|\lambda_k - \alpha_0| < |\lambda_k|$, the factored argument $|(\lambda_k - \alpha_0)\Delta x|$ remains smaller than the unfactored $|\lambda_k \Delta x|$. Consequently, as Δx decreases, the factored system reaches the small-argument regime, in which the leading-order comparison is asymptotically valid, at coarser streamwise resolutions than the unfactored march. This is consistent with the results reported here, in particular the randomised-forcing study of §5.3.4, where the factored system converges at substantially lower streamwise resolution than the unfactored march.

This argument requires the excited or physically relevant modes to cluster near one α_0 . A spectrum spread widely in λ cannot be served by a single carrier, and the unfactored march remains the fallback.

Resolution disc and the hypersonic spectrum example

At a fixed step the condition has a geometric reading. The unfactored system resolves a disc about the origin (Eq. (17)) and the factored system an equal disc about α_0 (Eq. (19)), so factoring translates the resolved disc to where the excited modes lie. Fig. 1 shows this for a Mach 4.5 boundary layer with $\alpha_0 = \omega$. At $\Delta x = C/|\omega|$ the factored disc covers the fast and slow discrete modes, while the unfactored disc covers only the fast.

Connection to PSE

The resolved-region picture is related to the PSE minimum step-size restriction. In PSE the downstream march is formally ill posed because the spectrum contains upstream-propagating, streamwise-elliptic modes.




Fig. 1 Eigenvalue spectrum for a Mach 4.5 boundary layer at frequency parameter $F = 220$ (see §5.3), with carrier wavenumber $\alpha_0 = \omega$. Background: $\tanh(\log(|\lambda - \omega|/|\lambda|))$, showing the leading-order per-mode resolution cost ratio (blue: carrier-wave cheaper; red: carrier-wave more expensive; black: break-even). Solid circle: factored region $|\lambda - \omega| \lesssim |\omega|$ at step size $\Delta x = C/|\omega|$. Dashed circle: unfactored region $|\lambda| \lesssim |\omega|$ at the same Δx .

The march remains stable only while these modes lie outside the resolved region, and refining Δx enlarges that region until they are captured, which fixes the minimum step [7,8,10]. Here the upstream-propagating modes are removed by the one-way projection (§3), so the resolved region stays a pure resolution concept and the PSE lower bound on Δx does not apply.

Scheme stability

Scheme stability is a separate question from resolution. A retained mode that lies in the integration scheme's unstable region is spuriously amplified, and because the carrier shift moves each mode's argument, factoring can move an otherwise stable mode into that region. This depends on the scheme and the spectrum and is not removed by the one-way projection. More generally, the modulus criterion Eq. (20) makes no distinction between phase and amplitude error, which need not be equally consequential in practice; the spurious amplification just described is the extreme case.

Slowly varying eigenfunctions

Parallel-flow analysis assumes the eigenfunctions v_k are independent of x . In non-parallel flows v_k varies with x , and the full computational benefit of the factoring additionally requires that this variation be resolved by the spatial discretisation:

$$\Delta x \left\| \frac{dv_k}{dx} \right\|_{\mathcal{H}} \ll \|v_k\|_{\mathcal{H}}. \quad (21)$$

This condition depends on the rapidity of base flow variation and the step size, not on α_0 . In typical boundary-layer flows the base flow evolves on scales much longer than the disturbance wavelength, so it is usually satisfied. In regions of strong non-parallelism (e.g. laminar separation bubbles or near the leading edge), finer Δx may be needed, but this requirement is a property of the base flow and applies equally to both factored and unfactored systems.

2.5 Closure condition

The wave-factored system Eq. (11) contains a free parameter: the carrier wavenumber α_0 . The resolution comparison in §2.4 shows that any computational benefit depends on α_0 lying near the excited spectral cluster: the question is how to choose α_0 in practice. Herbert [5] and Herbert [25] identify that different definitions of α used historically in PSE (for example the Rayleigh quotient, or the pointwise maximum of a disturbance

quantity) produce essentially the same results, because any α_0 within the cluster places the accuracy disc over the relevant eigenvalues. The apparent ambiguity provides additional freedom in the definition of α .

Iterating the Rayleigh quotient

The standard PSE closure condition [4,5] determines α_0 by iterating the Rayleigh quotient. Given an estimate $\alpha_0^{(n)}$, the factored system is solved for the envelope $\tilde{\phi}^{(n)}$ and a corrected carrier wavenumber is extracted as

$$\alpha_0^{(n+1)} = \alpha_0^{(n)} - i \frac{\left(\tilde{\phi}^{(n)}, \frac{d\tilde{\phi}^{(n)}}{dx} \right)_{\mathcal{H}}}{\|\tilde{\phi}^{(n)}\|_{\mathcal{H}}^2}, \quad n = 0, 1, \dots \quad (22)$$

Iteration continues until $|\alpha_0^{(n+1)} - \alpha_0^{(n)}| < \mathcal{E}$ for a prescribed non-dimensional tolerance $\mathcal{E}$. The same iterative scheme is used in adaptive LHNS [23,24].

For a disturbance dominated by a single eigenfunction the iteration converges to $\alpha_0 \approx \lambda_k$. The carrier wave tracks the dominant eigenvalue and the envelope reduces to the eigenfunction modulated by a slowly varying amplitude. This is the regime of maximum computational benefit. Convergence is typically achieved in 1–3 iterations for $\mathcal{E} = 10^{-9}$. The envelope can then easily be resolved at a coarse spatial grid spacing, since the residual wavenumber $|\lambda_k - \alpha_0|$ is small when $\alpha_0 \approx \lambda_k$ – the iteration identifies the α_0 that achieves this, rather than imposing slow variation on the envelope.

Behaviour under multi-modal disturbances

The iteration relies on the Rayleigh quotient providing a meaningful estimate of the dominant wavenumber. When several eigenfunctions carry comparable energy in a system whose eigenfunctions are approximately orthogonal under $\mathcal{H}$, the cross terms in the Rayleigh quotient are small and α_0 converges to a value within the convex hull of the excited eigenvalues (§2.2).

When the eigenfunctions are strongly non-orthogonal, cross terms contribute at leading order and the Rayleigh quotient can reside outside the convex hull of the excited eigenvalues. In this regime it can no longer be interpreted as a weighted-average wavenumber of the excited spectrum. Used as a closure condition, this has several consequences. The converged α_0 may lie outside the eigenvalue cluster, so that the accuracy region is no longer centred over the excited spectrum and any modes that require resolving may fall outside the translated region. Further, the iteration may converge to widely different values at adjacent streamwise stations, breaking smoothness of $\alpha_0(x)$. In severe cases it may fail to converge altogether.

Fixed carrier

An alternative that avoids these difficulties is to fix α_0 at a value known from the base flow. Several such choices cover the usual cases. All share advantages that the carrier is trivially smooth, requires no iterations, does not involve the Rayleigh quotient, and leaves the per-step cost identical to unfactored OWNS.

The carrier-wave should lie near the spectral content that dominates the physics of interest, so that eigenvalues of the excited modes fall well inside the accuracy disc. The trade-off is precision in the single-mode regime: the dominant eigenvalue lies close to α_0 but not exactly at it, so the envelope retains a slowly varying residual phase $\exp(i(\lambda_k - \alpha_0)x)$ that the iterated variant would remove. Since $|\lambda_k - \alpha_0| \ll |\alpha_0|$ for modes in the targeted spectral region, this residual is well within the resolution of any reasonable discretisation. This is demonstrated in Fig. 8, where the fixed-carrier calculation converges more slowly than the iterated system but faster than the unfactored system.

For problems dominated by discrete modes and the vorticity/entropy continuous spectrum, we set $\alpha_0 = \alpha_1 = (\omega - \beta W_\infty)/U_\infty$ from Eq. (B.1), the vorticity/entropy branch endpoint, where U_∞ and W_∞ are the non-dimensional freestream streamwise and spanwise velocities. In two-dimensional settings ($W_\infty = 0$ or $\beta = 0$) this reduces to $\alpha_0 = \omega$, as used in the hypersonic calculations of §5.3.

When the acoustic continuous spectrum carries significant energy, such as in receptivity, resolvent or optimal disturbance calculations, α_1 may under-resolve acoustic content. Setting $\alpha_0 = d_2$ (Eq. B.3), the acoustic branch midpoint, centres resolution at the acoustic behaviour. When both the acoustic and vorticity/entropy continuous spectra carry significant energy, $\alpha_0 = (\alpha_1 + d_2)/2$ can balance the two.

Unfactored OWNS is the limiting case $\alpha_0 = 0$, with no spectral preference.

Summary

In summary, iterating the Rayleigh quotient is the natural choice when a single eigenfunction dominates, where it provides maximum resolution benefit. When several eigenfunctions carry comparable energy and the excited eigenfunctions are strongly non-orthogonal, the fixed carrier is preferred. It avoids the difficulties

associated with cross-term contamination of the Rayleigh quotient while still centring the accuracy disc over the physically critical spectrum. The choice of carrier is application-dependent: α_1 is suited when discrete modes govern the instability growth; d_2 is suited to acoustic-dominated problems; the midpoint of d_2 and α_1 can balance acoustic and vortical/entropic contributions. All results in this paper use the iterating variant; in the swept-cylinder case (§5.2) and two hypersonic test cases (Mode-S forcing §5.3.1, randomised inlet forcing §5.3.4), the fixed carrier $\alpha_0 = \alpha_1$ is additionally computed and compared against the iterating variant.

3 OWNS and M-OWNS

Equation (3) represents a two-way evolution equation, where the semi-discrete operator $\mathbf{M}$ accommodates eigenfunctions that propagate in both downstream and upstream directions. Through application of Briggs' criterion [26], the finite spectrum of $\mathbf{M}$ can be partitioned into two subsets:

$$\sigma(\mathbf{M}) = \begin{cases} \text{downstream-propagating eigenvalues} \\ \text{upstream-propagating eigenvalues} \end{cases} \quad (23)$$

Stable downstream marching of Eq. (3) requires removing the upstream-propagating modes from $\mathbf{M}$; we therefore seek a parabolised operator $\mathbf{M}^+$ whose finite spectrum satisfies

$$\sigma(\mathbf{M}^+) = \begin{cases} \text{downstream-propagating eigenvalues} \\ \text{zero eigenvalues.} \end{cases} \quad (24)$$

We call $\mathbf{M}^+$ parabolised if it satisfies Eq. (24). The zero eigenvalues lie in the nullspace of $\mathbf{M}$ and are unaffected by the projection: they contribute nothing to the spatial dynamics regardless of whether they are formally assigned to $\mathbf{M}^+$.

3.1 The projection operator

Following Towne [27], Towne et al. [28], we construct a projection operator $\mathbf{P}$ from the eigen-decomposition Eq. (6) such that $\mathbf{M}^+ = \mathbf{P}\mathbf{M}$ retains only the downstream-propagating eigenvalues (Rudel et al. [29] give the corresponding continuous theory). On the finite-eigenvalue subspace,

$$\mathbf{P} = \mathbf{V}\mathbf{E}\mathbf{U}, \quad (25)$$

where $\mathbf{E}$ is a diagonal selector,

$$(\mathbf{E})_j = \begin{cases} 1 & \text{if } (\Delta)_j \text{ is downstream-propagating,} \\ 0 & \text{otherwise.} \end{cases} \quad (26)$$

The one-way operator is then

$$\mathbf{M}^+ = \mathbf{V}i\mathbf{E}\Delta\mathbf{U}, \quad (27)$$

whose spectrum satisfies Eq. (24). Since $\mathbf{E}$ and Δ are both diagonal, $\mathbf{P}$ and $\mathbf{M}$ commute:

$$\mathbf{P}\mathbf{M} = \mathbf{V}i\mathbf{E}\Delta\mathbf{U} = \mathbf{V}i\Delta\mathbf{E}\mathbf{U} = \mathbf{M}\mathbf{P}. \quad (28)$$

The projection operator $\mathbf{P}$ is constructed from the first-order pencil $(\mathbf{B}, \mathbf{A})$ alone. The second-order $\mathbf{C}$ terms are retained in the marching discretisation via backward differences but do not enter the projection construction. Since $\mathbf{P}$ does not in general commute with $\mathbf{C}$, the upstream-downstream splitting of the full second-order system may in principle differ from that of the first-order pencil. This question has not, to the authors' knowledge, been addressed in the literature. The $\mathbf{C}$ terms contain only the second order streamwise viscous derivatives, which are small (scaling with the reciprocal of the Reynolds number) relative to $\mathbf{A}$ and $\mathbf{B}$; the results and conclusions of Zhu and Towne [13] and our own numerical tests across subsonic and hypersonic configurations show no evidence of error attributable to this omission.

Projection of the carrier-wave factored system We now show that the projection operator $\mathbf{P}$ parabolises the carrier-wave factored operator $\tilde{\mathbf{M}}$, provided it parabolises $\mathbf{M}$. Since $\mathbf{P}$ parabolises $\mathbf{M}$, the matrix $\mathbf{E}$ satisfies Eq. (26). Expanding the factored operator $\tilde{\mathbf{M}}$ from Eq. (12) gives

$$\tilde{\mathbf{M}} = \mathbf{V} (i\Delta - i\alpha_0 \mathbb{I}) \mathbf{U}. \quad (29)$$

Applying $\mathbf{P}$ and using the biorthogonality condition Eq. (5) yields

$$\mathbf{P}\tilde{\mathbf{M}} = \mathbf{V}\mathbf{E}\mathbf{U}\mathbf{V} (i\Delta - i\alpha_0 \mathbb{I}) \mathbf{U} = \mathbf{V} i\mathbf{E} (\Delta - \alpha_0 \mathbb{I}) \mathbf{U}. \quad (30)$$

From Eq. (26), the operator $\mathbf{E} (\Delta - \alpha_0 \mathbb{I})$ has the property

$$[\mathbf{E} (\Delta - \alpha_0 \mathbb{I})]_j = \begin{cases} (\Delta - \alpha_0 \mathbb{I})_j & \text{if } (\Delta)_j \text{ is downstream-propagating,} \\ 0 & \text{otherwise,} \end{cases} \quad (31)$$

since $\mathbf{E}$ selects the same entries regardless of the shift by α_0 . Therefore $\tilde{\mathbf{M}}^+ = \mathbf{P}\tilde{\mathbf{M}}$ is a one-way operator: the same projection that parabolises the unfactored system also parabolises the carrier-wave factored system. No modification of $\mathbf{P}$ or the recursion parameter selection is required.

3.2 The approximate projection operator: OWNS-R

The eigenvalue decomposition in the projection operator $\mathbf{P}$, Eq. (25), is prohibitively expensive to evaluate at each spatial marching step. Zhu and Towne [13] proposed the OWNS-R technique, which approximates $\mathbf{P}$ via an operator $\mathbf{P}_{N_\beta}$ that depends on a sequence of complex recursion parameters

$$\beta^\pm = (\beta_j^+, \beta_j^-)_{j=1, \dots, N_\beta} \subset \mathbb{C}^2, \quad N_\beta \in \mathbb{N}. \quad (32)$$

The sets β^+ and β^- are called the positive and negative recursion parameters respectively, and are placed near the downstream- and upstream-propagating eigenvalues of $\mathbf{M}$. With appropriately placed $\beta^\pm$, the approximation satisfies $\lim_{N_\beta \rightarrow \infty} \mathbf{P}_{N_\beta} = \mathbf{P}$.

One form of the approximate projection operator is

$$\mathbf{P}_{N_\beta} = \left(\mathbb{I} + c \prod_{k=1}^{N_\beta} (\mathbf{M} - i\beta_k^+ \mathbb{I}) (\mathbf{M} - i\beta_k^- \mathbb{I})^{-1} \right)^{-1}, \quad (33)$$

where $c > 0$ is a tuning parameter that balances retention of downstream modes against removal of upstream modes; typical values are $c = 0.1$ – 10 . Clearing the inner inverses and refactoring yields

$$\mathbf{P}_{N_\beta} = \frac{1}{c+1} \prod_{k=1}^{N_\beta} (\mathbf{M} - i\beta_k^* \mathbb{I})^{-1} (\mathbf{M} - i\beta_k^- \mathbb{I}), \quad (34)$$

where $\beta_j^* \in \mathbb{C}$ are the roots of the N_β -th order polynomial

$$c \prod_{k=1}^{N_\beta} (\lambda - \beta_k^+) + \prod_{k=1}^{N_\beta} (\lambda - \beta_k^-) = (c+1) \prod_{k=1}^{N_\beta} (\lambda - \beta_k^*). \quad (35)$$

In practice, $\mathbf{M}$ is never formed; all operations are expressed in terms of the pencil $(\mathbf{B}, \mathbf{A})$. Equation (34) takes the form

$$\mathbf{P}_{N_\beta} = \frac{1}{c+1} \prod_{k=1}^{N_\beta} (\mathbf{B} - i\beta_k^* \mathbf{A})^{-1} (\mathbf{B} - i\beta_k^- \mathbf{A}), \quad (36)$$

Table 1 Number of linear solves per streamwise step for each method. N_{it} denotes the number of Rayleigh quotient iterations (typically 2–3 for $\mathcal{E} = 10^{-9}$) and N_β denotes the number of recursion parameters (typically 20–60).

Method	Solves per step	Typical range
PSE	$1 + N_{it}$	3–4
OWNS	$1 + N_\beta$	21–61
M-OWNS (iterating)	$1 + N_{it} + N_\beta$	23–64
M-OWNS (non-iterating)	$1 + N_\beta$	21–61

and the action of $\mathbf{P}_{N_\beta}$ on a state vector ϕ is computed via the recursion

$$\phi_0 = \frac{1}{c+1}\phi, \quad (37)$$

$$(\mathbf{B} - i\beta_k^* \mathbf{A}) \phi_k = (\mathbf{B} - i\beta_k^- \mathbf{A}) \phi_{k-1}, \quad k = 1, \dots, N_\beta, \quad (38)$$

$$\phi^+ = \phi_{N_\beta}. \quad (39)$$

Since each solve involves the shifted pencil $(\mathbf{B} - i\beta_k^* \mathbf{A})$, and the β_k^* are finite and distinct from the eigenvalues of $(\mathbf{B}, \mathbf{A})$, the infinite eigenvalues associated with singular $\mathbf{A}$ do not enter the recursion and require no special treatment.

Each recursion step requires one banded linear solve of order $5n_y$, so the cost of applying $\mathbf{P}_{N_\beta}$ is N_β solves.

3.3 The M-OWNS marching algorithm

We now describe the complete marching algorithm. Since $\mathbf{P}$ and $\mathbf{M}$ commute, the two-way system can be solved first and the result projected, rather than forming and marching with $\mathbf{M}^+$ directly. The carrier-wave factored system Eq. (14) provides the most general framework; unfactored OWNS is recovered by setting $\alpha_0 = 0$.

At each streamwise station, the algorithm proceeds as follows:

1. **Solve.** Advance the carrier-wave factored two-way system one step using backward differencing to obtain $\tilde{\phi}$.
2. **Iterate** (optional). Update the carrier wavenumber α_0 via the Rayleigh quotient Eq. (22) and re-solve step 1 with the updated α_0 . Repeat until $|\alpha_0^{(n+1)} - \alpha_0^{(n)}| < \mathcal{E}$.
3. **Project.** Apply the approximate projection $\tilde{\phi}^+ = \mathbf{P}_{N_\beta} \tilde{\phi}$ using the recursion Eqs. (37)–(39).

Three cases arise as parameter choices within this single framework:

- *Iterated α_0* (iterated M-OWNS): the carrier wavenumber is updated at each station via the Rayleigh quotient.
- *Fixed $\alpha_0 = \alpha_1$* (non-iterating M-OWNS): the carrier wavenumber is set to the vorticity/entropy branch endpoint $(\omega - \beta W_\infty)/U_\infty$ from Eq. (B.1), which is known *a priori*. No iterations are required.
- $\alpha_0 = 0$ (unfactored OWNS): the carrier-wave factoring is absent and the unfactored two-way system is solved directly. No iterations are required.

Computational cost At each streamwise station, the marching solve (one banded linear system of order $5n_y$) is performed $(N_{it} + 1)$ times: the initial solve plus N_{it} Rayleigh quotient iterations. A single projection requiring N_β recursion solves follows. The iteration and projection are sequential, not nested: each iteration repeats only the marching solve, and the projection is applied once after the iteration has converged. Table 1 summarises the per-step cost for each method.

Non-iterating M-OWNS has identical per-step cost to unfactored OWNS. The iteration cost N_{it} is additive at each step, not multiplicative on the projection cost N_β ; the iterating variant therefore incurs an overhead of $N_{it}/(1 + N_\beta)$ relative to OWNS, which is under 15% across the typical parameter range. PSE is cheapest per step but is subject to a minimum step-size restriction that limits its applicability [10]. The total cost additionally depends on the number of streamwise steps n_x required for a given accuracy; quantitative comparisons are given in §5.4.

Projection robustness The coarser step size permitted by carrier-wave factoring also eases the demands on the approximate projection. At a small step size the integrator’s spectral window is wide, so $\mathbf{P}_{N_\beta}$ must suppress

a broad portion of the upstream-propagating spectrum. Since the OWNS-R recursion has a practical upper bound on N_β (§3.4), a wide spectral window requires more carefully positioned recursion parameters. M-OWNS resolves the cluster of physically relevant modes near α_1 at a larger step size, narrowing the spectral window so that fewer upstream modes fall within it and the same N_β provides a more faithful projection. By the same leading-order argument, M-OWNS should not require a better projection than unfactored OWNS for the same physical accuracy, provided α_0 is never closer to the upstream spectrum than the origin.

3.4 Recursion parameter selection

The recursion parameters $\beta^\pm$ introduced in §3.2 must be chosen so that $\mathbf{P}_{N_\beta}$ faithfully approximates the exact projection $\mathbf{P}$: the positive parameters β^+ are placed near downstream-propagating eigenvalues to ensure their retention, and the negative parameters β^- are placed near upstream-propagating eigenvalues to ensure their removal. The tuning parameter c balances the relative strength of retention and removal.

The general strategy follows the geometric framework of Towne and Colonius [11]. Each recursion parameter pair (β_k^+, β_k^-) defines a rational factor that drives the scalar projection function $\mathcal{P}_{N_\beta}$ (the scalar version of Eq. (33)) toward one near β_k^+ and toward zero near β_k^- , with exponential convergence as N_β increases. This convergence is limited in practice by the accumulation of numerical error in the recursive inversions, as first observed by Sleeman and Colonius [14]. The parameter placement must reflect the spectral structure of the flow, which varies with Mach number, spanwise wavenumber and streamwise location. The wall-normal grid stretching produces a non-uniform sampling of the continuous spectral branches, with eigenvalues concentrated near the branch endpoints and increasingly sparse further along each branch. The recursion parameters must follow this distribution: clustered near the branch endpoints and more widely spaced away from them.

Sleeman and Colonius [14] proposed a greedy algorithm for the automatic selection of recursion parameters. Their method uses knowledge of the spectrum and applies the greedy selection to determine which eigenvalues to use as recursion parameters.

The parameters are recalculated at each streamwise station as the base flow evolves and the spectrum migrates. This is standard practice in OWNS implementations. The explicit parameter expressions for the subsonic and supersonic configurations studied in this work, together with their derivation and validation, are given in Appendix B.

4 Numerical methods

The stability calculations are performed using an in-house compressible spatial marching code incorporating LST and PSE [17]. The OWNS-R and M-OWNS results reported here are from an independently developed solver.

All spatial quantities are non-dimensionalised by δ_0 , the boundary-layer thickness at the inlet $\bar{x}_0$. The local Reynolds number based on boundary-layer thickness is $R_\delta = \sqrt{\bar{x}} R_\infty$, where R_∞ is the freestream Reynolds number and $\bar{x}$ is the dimensional distance from the origin. The inlet Reynolds number is $R = R_\delta(\bar{x}_0)$. The base flow is non-dimensionalised by freestream values at the inlet.

The non-dimensional frequency and spanwise wavenumber used throughout this work are

$$F = \frac{\omega \times 10^6}{R}, \quad b = \frac{\beta \times 10^3}{R}, \quad (40)$$

where ω and β are the non-dimensional temporal frequency and spanwise wavenumber introduced in Eq. (1).

The wall-normal domain $0 < \bar{y} < L_{\bar{y}}$ is discretised with n_y points using fourth-order finite differences on the stretched grid of Malik [30]:

$$y_{\text{Malik}}(\eta, y_{\text{max}}, y_{\text{half}}) = \frac{a \eta}{b - \eta}, \quad (41)$$

where $0 \leq \eta \leq 1$ is the uniform computational coordinate, y_{max} is the far-field truncation height, and more than half the grid points lie between 0 and y_{half} , with

$$a = \frac{y_{\text{max}} y_{\text{half}}}{y_{\text{max}} - 2 y_{\text{half}}}, \quad b = 1 + \frac{a}{y_{\text{max}}}. \quad (42)$$




Fig. 2 (a) Rayleigh quotient and (b) growth-rate comparisons for a 2D-disturbance ($b = 0$) with $F = 86$ in a subsonic boundary-layer with variable Mach number. (—) $M = 0.02$, (—) $M = 0.2$, (—) $M = 0.5$, (—) $M = 0.8$. All calculations include PSE with $n_x = 110$, M-OWNS with $n_x = 110$ and $n_x = 1000$, and OWNS at $n_x = 1000$.

This stretching function is also used to distribute recursion parameters non-uniformly in the complex plane (Appendix B).

The streamwise domain $0 < \bar{x} < L_{\bar{x}}$ is discretised with n_x points using second-order backward differencing (BFD2), unless stated otherwise. Higher-order backward-difference schemes have smaller stability regions for stiff systems and are not pursued here. Grid extents and resolutions are specified for each flow configuration in the relevant subsection.

At the wall, no-slip conditions are imposed on velocity disturbances with either a zero-disturbance or adiabatic condition for temperature. In the freestream, Thompson characteristic conditions suppress numerical reflections [31]. Initial conditions are provided by one of three methods: a locally parallel eigenfunction from LST, a summation of different freestream waves, or random noise. Disturbances can also be generated via wall suction and blowing. The specific choice is stated for each case.

For iterated M-OWNS and PSE, the convergence criterion $\mathcal{E}$, non-dimensionalised by the boundary-layer thickness at the inlet, denotes the absolute difference between successive wavenumbers at each station. We use $\mathcal{E} = 10^{-9}$ unless otherwise stated, yielding $N_{it} = 2\text{--}3$ iterations per station.

5 Results

We assess M-OWNS across three flow configurations: subsonic flat-plate boundary-layers (§5.1), incompressible flow over a swept cylinder (§5.2), and a Mach 4.5 hypersonic boundary-layer (§5.3). In every case, the spectral resolution comparison (§2.4) correctly predicts the resolution requirements of the two methods.

5.1 Subsonic Flat-Plate boundary layers

We consider isothermal flat-plate flows at freestream Reynolds number $R_{\infty} = 10^6 \text{ m}^{-1}$ and temperature $T_{\infty} = 298 \text{ K}$, with Mach numbers $M = 0.02, 0.2, 0.5$ and 0.8 . The computational domain is initialised at $R_{\delta} = 400$ using a locally parallel eigenfunction (LST) corresponding to the most unstable T-S mode. The analysis considers frequency $F = 86$ and spanwise wavenumbers $b = 0, 0.1, 0.2$ and 0.3 . All calculations employ $n_y = 221$ wall-normal points. For this flat-plate geometry the curvature $\kappa = 0$, so the metric term $\chi = 1$ in Eq. (A.1). For every two-dimensional subsonic case, we use $N_{\beta} = 30$ recursion parameters; exact parameters are given in §B.2.

5.1.1 Validation across Mach number and obliqueness

Figure 2 presents Rayleigh quotient and growth-rate comparisons across Mach number for a two-dimensional disturbance ($b = 0$). M-OWNS at both low resolution ($n_x = 110$) and high resolution ($n_x = 1000$) agrees with PSE ($n_x = 110$) and OWNS ($n_x = 1000$) throughout the domain. Figure 3 demonstrates equivalent agreement across spanwise wavenumbers $b = 0\text{--}0.3$ in an incompressible boundary-layer ($M = 0.02$), confirming that M-OWNS accurately models three-dimensional oblique disturbances.

Downstream, the Rayleigh quotient transitions from the T-S value to that of the vorticity wave, $\alpha_1 = \omega$ in Eq. (B.1). Non-parallel base flow evolution transfers energy from the decaying T-S wave to the more slowly decaying vorticity mode, which eventually dominates. This transition is resolved by OWNS and M-OWNS at sufficient n_x .




Fig. 3 (a) Rayleigh quotient and (b) growth-rate comparisons for a 3D-disturbance with $F = 86$ and variable spanwise wavenumber in an incompressible boundary-layer ($M = 0.02$). (—) $b = 0$, (—) $b = 0.1$, (—) $b = 0.2$, (—) $b = 0.3$. All calculations include PSE with $n_x = 110$, M-OWNS with $n_x = 110$ and $n_x = 1000$, and OWNS at $n_x = 1000$.




Fig. 4 (Top) Rayleigh quotient and (bottom) growth-rate comparisons for a disturbance with $F = 86$ in an incompressible boundary-layer ($M = 0.02$) with zero spanwise wavenumber ($b = 0$). Only M-OWNS gives consistent and converged results with $n_x = 200, 1000$; both OWNS and PSE with removed streamwise pressure gradient ($p_x = 0$) fail on using $n_x = 200$ points, while PSE with p_x retained has the worst behaviour.

5.1.2 M-OWNS in the intermediate resolution regime

Figure 4 demonstrates a critical advantage of M-OWNS at intermediate streamwise resolution. At $n_x = 200$, both PSE and unfactored OWNS fail: PSE violates its step-size constraint, while OWNS exhibits oscillations leading to breakdown. M-OWNS maintains numerical stability and produces results consistent with the converged high-resolution calculations. This intermediate regime, too coarse for OWNS and too fine for PSE, is precisely where the resolution benefit predicted by the spectral resolution comparison (§2.4) is most consequential.

5.1.3 Convergence of disturbance amplitude and energy

While normalised quantities such as Rayleigh quotient require relatively modest resolution, accurate prediction of disturbance amplitude and energy demands considerably finer x -discretisation with unfactored OWNS. Fig. 5 shows that OWNS requires $n_x > 8000$ for converged maximum velocity and kinetic energy in an incompressible two-dimensional boundary-layer with $F = 86$. M-OWNS achieves converged predictions with $n_x \geq 150$.

Similar resolution requirements emerge in compressible flows. Fig. 6 compares energy predictions for a Mach 0.8 boundary-layer with $b = 0.1$ and $F = 150$, using implicit Euler (BFD1) and BFD2 discretisations with $\mathcal{E} = 10^{-9}$. OWNS energy predictions increase gradually with n_x , achieving agreement with M-OWNS only at extreme discretisations ($n_x > 4000$). First-order differencing performs worst, as expected. M-OWNS demonstrates convergent behaviour across the entire resolution range.




Fig. 5 (Top) Maximum streamwise velocity magnitude and (bottom) disturbance energy for a disturbance with $F = 86$ in an incompressible 2D boundary-layer ($M = 0.02$) with zero spanwise wavenumber ($b = 0$).




Fig. 6 Maximum density-scaled kinetic energy ($E_{\mathcal{H}}$) versus streamwise resolution (n_x) for a $M = 0.8$ flat-plate boundary-layer, with $b = 0.1$ and frequency $F = 150$.

5.2 Swept cylinder

To test M-OWNS in a three-dimensional flow with evolving spectra, we consider the effectively incompressible ($M = 0.0001$) boundary layer over a swept cylinder with radius 0.5 m and sweep angle 65° , following Malik et al. [16]. We set $T_\infty = 300$ K and with $n_y = 301$ wall-normal points over the domain $132 \leq R_\delta \leq 1460$ and $N_\beta = 20$ recursion parameters. The curvature $\bar{\kappa} = 2/\text{m}$ is included via Eq. (A.1). Two crossflow disturbances are examined at $\beta = 0.4$ ($b = 3.030$): a stationary mode ($F = 0$) and a travelling mode ($F = 75$).

This flow introduces a computational challenge absent in flat-plate cases – the vorticity/entropy branches migrate toward the acoustic spectrum as the crossflow ratio W_∞/U_∞ decreases downstream (Fig. 25). The recursion parameters must track this evolution; details of the adaptive strategy are given in Appendix B.

Figure 7 demonstrates agreement between PSE, M-OWNS, and OWNS-R for both stationary and travelling crossflow modes. Fig. 8 compares convergence with streamwise resolution for three formulations: standard OWNS, M-OWNS with the iterated Rayleigh quotient ($\alpha_0 = \alpha$), and M-OWNS with the fixed carrier $\alpha_0 = \alpha_1 = -\beta W_\infty/U_\infty$.

The ordering reflects spectral distance. At the inlet, the crossflow wave and the vorticity/entropy branches are far away from the origin/acoustic branches. As the flow develops downstream, the crossflow W_∞/U_∞ and the crossflow wave and the vorticity/entropy branches move towards the origin. The iterated carrier sits on and tracks the crossflow eigenvalue at every station, so the residual wavenumber is zero and convergence requires only $n_x \sim 150$. The fixed carrier α_1 mode lies close to but not on the eigenvalue, leaving a nonzero




Fig. 7 (a) Rayleigh quotient and (b) growth-rate comparisons for crossflow disturbances ($\beta = 0.4$) in the swept cylinder analysis. Results shown for (—) travelling ($F = 75$) and (—) stationary ($F = 0$) crossflows. Computations performed using PSE ($n_x = 150$), M-OWNS ($n_x = 150, 2000$), and OWNS ($n_x = 2000$).




Fig. 8 Maximum streamwise velocity for stationary crossflow disturbances ($\beta = 0.4$, $F = 0$) in the swept cylinder analysis. Top: OWNS; middle: M-OWNS with iterated $\alpha_0 = \alpha$; bottom: M-OWNS with fixed carrier-wave with $\alpha_0 = \alpha_1 = -\beta W_\infty / U_\infty$.

residual wavenumber throughout the march; convergence requires $n_x \sim 1000$. Standard OWNS measures all wavenumbers from the origin, the largest spectral distance, and requires $n_x \sim 8000$. The three panels therefore illustrate the resolution criterion of §2.4 in quantitative terms: convergence rate is governed by the distance from the carrier to the dominant eigenvalue, not by the absolute wavenumber of the mode.

5.3 Hypersonic boundary-layer

We examine a Mach 4.5 boundary layer following Mack [1], with $R_\infty = 7.2 \times 10^6 \text{ m}^{-1}$, $T_\infty = 65.15 \text{ K}$, $p_\infty = 728.44 \text{ Pa}$, Prandtl number $\sigma = 0.72$, frequency $F = 220$ and $b = 0$. The adiabatic wall condition $T_y = 0$ is imposed. This configuration supports multiple discrete modes: the slow mode (S) and fast mode (F) of Fedorov [32], which synchronise far downstream and generate acoustic radiation [33]. We denote the associated continuous spectral branch end-points as slow acoustic (SA), fast acoustic (FA), and vorticity (V).




Fig. 9 Wall pressure amplitude $|p_{\text{wall}}|$ for the Mach 4.5 boundary layer at $F = 220$. Top: OWNS; bottom: M-OWNS. DNS reference from Ma and Zhong [34]. OWNS converges between $n_x = 4000$ and 8000 ; M-OWNS is converged at all resolutions tested.

Four forcing scenarios probe different aspects of M-OWNS capability: (i) Mode S inlet forcing to study synchronisation, (ii) wall suction/blowing for multimodal excitation, (iii) simultaneous freestream wave forcing with SA, FA and V waves to test the broadband, disparate-eigenvalue case, and (iv) randomised inlet forcing, where the absence of a dominant upstream mode provides a direct test of the leading-order spectral resolution condition and a comparison of the iterated and fixed-carrier closures.

Supersonic parameter placement is distinct from subsonic and is discussed in §B.3. All calculations use a BFD2 scheme.

5.3.1 Mode S forcing

The computational domain spans $R_\delta = 400$ to 1410 with $n_y = 600$ and $N_\beta = 40$. The inlet is forced with a Mode S eigenfunction from LST.

The wall pressure $|p_{\text{wall}}|$ tracks the modal content at the surface. Fig. 9 shows that OWNS converges between $n_x = 4000$ and 8000 ; M-OWNS is converged at all resolutions tested, down to $n_x = 650$.

The pressure field $\text{Re}(p)$ reveals the acoustic radiation generated by S–F synchronisation in the freestream. Fig. 10 compares OWNS and M-OWNS at $n_x = 12\,000$ and 2000 . At the finer resolution, both methods capture the near-wall second mode and its freestream radiation. At $n_x = 2000$, OWNS fails to resolve the radiation pattern while M-OWNS maintains both near-wall accuracy and radiation field structure. At $n_x = 1000$ (not shown), M-OWNS still resolves the near wall Mode S growth but not the freestream radiation; OWNS captures neither.

Non-iterating M-OWNS

We repeat the Mode S calculation with the fixed carrier $\alpha_0 = \alpha_1 = \omega$ (§2.5), so that $N_{\text{it}} = 0$ and the per-step cost is identical to OWNS. Fig. 11 compares iterated M-OWNS ($\mathcal{E} = 10^{-14}$), non-iterating M-OWNS, and OWNS, all at $n_x = 1000$. The two M-OWNS variants are indistinguishable in both Rayleigh quotient and wall pressure. OWNS at the same n_x under-predicts the amplitude by more than an order of magnitude.

This is a comparison at equal per-step cost and equal n_x . For all resolutions tested, the iterated and non-iterating variants produce identical results. Non-iterating M-OWNS captures the S–F synchronisation and second-mode amplification that OWNS at the same resolution does not. As predicted by the spectral resolution comparison, the physically relevant eigenvalues cluster near α_1 , so the fixed carrier centres the accuracy disc over the excited spectrum without iteration.




Fig. 10 Pressure field $\text{Re}(p)$ for Mode S forcing in the Mach 4.5 boundary layer. Left: OWNS; right: M-OWNS. Top row: $n_x = 12\,000$; bottom row: $n_x = 2\,000$. At the coarser resolution, OWNS fails to resolve the acoustic radiation field that M-OWNS captures. The dashed white line denotes the boundary-layer edge.




Fig. 11 Left: Rayleigh quotient; right: wall pressure, for Mode S forcing in the Mach 4.5 boundary layer at $n_x = 1\,000$. Three methods are compared: iterated M-OWNS ($\mathcal{E} = 10^{-14}$), non-iterating M-OWNS ($\alpha_0 = \omega$), and OWNS. In the left panel the non-iterating curve shows the fixed carrier $\alpha_0 = \omega$ rather than a Rayleigh quotient; the Rayleigh quotient recovered from the non-iterating solution (not shown) is identical to the iterated one. The non-iterating variant has identical per-step cost to OWNS.

5.3.2 Wall forcing via suction and blowing

The wall forcing of Ma and Zhong [35,36] is applied:

$$v_{\text{wall}}(\bar{x}) = \begin{cases} a_0 \sin(2\pi(\bar{x} - \bar{x}_0)/\bar{x}_l), & \text{if } |\bar{x} - \bar{x}_0| \leq \bar{x}_l/2, \\ 0, & \text{otherwise,} \end{cases} \quad (43)$$

where $\bar{x}_0 = 0.0125$ m ($R_\delta = 300$), $\bar{x}_l = 0.0029127$ m, and a_0 is arbitrary. The forcing extends over $282 \leq R_\delta \leq 317$. The spatial scales of the forcing are smaller than PSE's minimum stable step size, precluding its use. The domain is $275 \leq R_\delta \leq 1275$ and $n_y = 721$ points in the wall-normal direction, with $N_\beta = 40$.

Figure 12 shows the normalised wall pressure. M-OWNS at $n_x = 1\,000$ agrees with LHNS at $n_x = 8949$ throughout the domain. OWNS at $n_x = 1\,000$ under-predicts the amplitude. The pressure and temperature fields (Fig. 13) confirm that M-OWNS and LHNS produce nearly identical distributions. LHNS is used as the reference here because this is the only case where upstream-propagating content could differ materially between the one-way and elliptic treatments. Minor differences in the pressure field near the forcing region are attributable to the projection operator removing upstream-propagating content in the subsonic sublayer, which the LHNS retains.

The normalised wall-normal profiles at $R_\delta \approx 1\,050$ (Fig. 14) show agreement in eigenfunction shape across all three methods. OWNS captures the shape correctly but not the streamwise amplitude.




Fig. 12 Normalised wall pressure $|p_{\text{wall}}|$ for the $M = 4.5$ boundary layer forced by surface suction/blowing (Eq. 43). The dashed black curve is M-OWNS forced with a Mode S eigenfunction at the inlet, shown for reference.




Fig. 13 Disturbance pressure $\text{Re}(p)$ (left) and temperature $\text{Re}(T)$ (right) for the $M = 4.5$ boundary layer with surface suction/blowing. Top: OWNS, $n_x = 1000$; middle: M-OWNS, $n_x = 1000$; bottom: LHNS, $n_x = 8949$. The white dashed line denotes the boundary-layer edge.




Fig. 14 Wall-normal profiles of temperature $|T|$ (left) and pressure $|p|$ (right) at $R_\delta = 1050.6$. M-OWNS and OWNS at $n_x = 1000$; LHNS at $n_x = 8949$. Temperature is normalised by its maximum. All three methods agree in eigenfunction shape; the unnormalised maximum temperature for OWNS is 0.0241, compared with 0.4981 for M-OWNS, a factor of 20. The slight offset at $y^* = 0$ is due to the effect of the one-way projection. Dashed line: boundary-layer edge.




Fig. 15 Rayleigh quotient (left) and wall pressure (right) for individual freestream wave forcing in the Mach 4.5 boundary layer ($F = 220$, $b = 0$): Mode S (—), Mode F (—), SA wave (---), FA wave (---), and V wave (---). Each wave is computed individually with M-OWNS at $n_x = 1000$. These evolutions serve as reference for the combined forcing in Fig. 16.

5.3.3 Multi-Mode freestream forcing

The preceding test cases produced responses dominated by Mode S. To probe M-OWNS performance when multiple spectrally distinct wave types contribute simultaneously, we initialise the computation with a superposition of three freestream wave types: slow acoustic (SA, $\alpha = \omega M / (M - 1)$), fast acoustic (FA, $\alpha = \omega M / (M + 1)$), and vorticity (V, $\alpha = \omega$), each with equal magnitude and the total disturbance energy normalised to unity. No LST eigenfunctions are used. The computation uses $n_y = 721$, $N_\beta = 60$ and marches from $R_\delta = 270$ to 1321. This configuration simultaneously excites modes spanning a wide range of Rayleigh quotients, as shown by the individual wave evolutions in Fig. 15.

Despite this spread, the spectral resolution comparison (§2.4) predicts that M-OWNS should retain its resolution advantage. The physically relevant criterion is not eigenvalue proximity to one another, but proximity to the carrier α_0 . All three wave types are excited at the forcing frequency ω , so their eigenvalues cluster near ω .

Figure 16 confirms this prediction. OWNS at $n_x = 1000$ exhibits large Rayleigh quotient oscillations and wall pressure that departs from the converged solution. M-OWNS at $n_x = 1000$ is indistinguishable from the $n_x = 8000$ reference in both quantities. The pressure fields (Fig. 17) show the same result: OWNS at $n_x = 1000$ produces unphysical structure absent from the $n_x = 8000$ reference; M-OWNS at $n_x = 1000$ agrees with the reference throughout.

5.3.4 Randomised inlet forcing

The multi-mode forcing of §5.3.3 prescribes specific wave types at the inlet. To test robustness under less controlled conditions, we replace the structured inlet with a randomised disturbance profile. Each component of the state vector at every wall-normal grid point is multiplied by an independent random factor drawn uniformly from $[-1, 1]$. The profile is truncated to zero above $y^* = 10\text{mm}$ to preserve the freestream




Fig. 16 Multi-mode freestream forcing (SA + FA + V) in the Mach 4.5 boundary layer. Left: OWNS; right: M-OWNS. Top row: Rayleigh quotient; bottom row: wall pressure. Resolutions: $n_x = 1000, 2000, 4000$ and 8000 . The individual SA, FA and V Rayleigh quotients from Fig. 15 are shown for reference. OWNS exhibits strong resolution dependence at $n_x \leq 2000$; M-OWNS is converged at all resolutions.

boundary conditions (Fig. 18). The computation uses $n_y = 721$, $N_\beta = 60$ and marches from $R_\delta = 270$ to 1321. Despite the incoherent inlet, the field rapidly organises into a coherent wave structure (Fig. 21).

Figure 19 compares wall pressure convergence for OWNS, iterated M-OWNS ($\alpha_0 = \alpha$), and fixed-carrier M-OWNS ($\alpha_0 = \omega$). At $n_x = 1000$, OWNS and M-OWNS resolve complementary parts of the spectrum. In the upstream region ($R_\delta \lesssim 500$), OWNS captures small-scale oscillations in $|p_{\text{wall}}|$ but not the correct amplitude, while M-OWNS captures the amplitude but not the oscillations. This is consistent with the spectral resolution condition (§2.4): eigenvalues nearer the origin than ω satisfy $|\lambda_k| < |\lambda_k - \omega|$ and are better resolved by the unfactored system, while eigenvalues near ω are better resolved by the factored system. The Rayleigh quotient data (Fig. 20) support this: α_r lies below ω in the upstream region, indicating that the dominant content has wavenumber closer to the origin than to ω .

Downstream, Mode S growth through S–F synchronisation produces content with Rayleigh quotient near ω . OWNS at $n_x = 1000$ loses the solution by $R_\delta \approx 600$; $n_x = 3000$ begins to under-resolve Mode S growth by $R_\delta \approx 700$; convergence requires $n_x \approx 6000$. Both M-OWNS variants resolve the full domain by $n_x = 3000$.

The fixed carrier $\alpha_0 = \omega$ performs comparably to the iterated Rayleigh quotient throughout. The iteration converges for this realisation at every streamwise step, despite the Rayleigh quotient being used as a recursion parameter (see §B.3). The fixed carrier requires no convergence and so avoids the cross-term sensitivity of the Rayleigh quotient under strongly non-orthogonal excitation (§2.2); this non-orthogonality may cause the iterations to fail for different initial conditions (§2.5).

5.4 Total cost

The total cost of each method scales as n_x times the per-step cost of Table 1. To leading order, the step count n_x required for a given accuracy is set by the spectral resolution comparison of §2.4: carrier-wave factoring reduces the residual wavenumbers of modes near α_0 , permitting coarser step sizes for any consistent spatial discretisation. Beyond leading order the gain depends on the scheme and on the eigenvalue phase; the convergence studies below report the n_x each method requires in practice.

The randomised-inlet case of §5.3.4 is the most stringent test of M-OWNS. Both M-OWNS variants converge at $n_x = 3000$ while OWNS requires $n_x = 6000$. With $N_\beta = 60$, this corresponds to 3.7×10^5 solves




Fig. 17 Pressure field $\text{Re}(p)$ for multi-mode freestream forcing (SA + FA + V) in the Mach 4.5 boundary layer. From top to bottom: OWNS at $n_x = 1000$, M-OWNS at $n_x = 1000$, OWNS at $n_x = 8000$, M-OWNS at $n_x = 8000$. White dashed line: boundary-layer edge; white solid line: sonic line. At $n_x = 1000$, OWNS fails to resolve the near-wall second-mode structure that M-OWNS captures; at $n_x = 8000$ both methods agree.




Fig. 18 Randomised initial condition at $R_\delta = 270$. The profile is set to zero above $y^* = 10$ mm. Dashed line: boundary-layer edge.

for OWNS. For iterating M-OWNS with $N_{\text{it}} = 3$, the total solve count is 1.9×10^5 , a reduction of 47.5%. For fixed-carrier M-OWNS the total solve count is 1.8×10^5 , a reduction of 50.0%. Even larger reductions are obtained in less stringent cases.

The swept-cylinder configuration of §5.2 uses $N_\beta = 20$ and requires $n_x = 8000$ for OWNS, corresponding to 1.7×10^5 solves. Iterating M-OWNS converges at $n_x = 150$ with $N_{\text{it}} = 3$, giving 3.6×10^3 solves, a factor of 47 below OWNS. Fixed-carrier M-OWNS converges at $n_x = 1000$, giving 2.1×10^4 solves, a factor of 8 below OWNS. OWNS resolves the full inter-branch span with $\alpha_0 = 0$; both M-OWNS variants centre the accuracy disc near the crossflow wavenumber, with the iterating variant tracking it exactly.




Fig. 19 Wall pressure $|p_{\text{wall}}|$ for randomised inlet forcing. Top: OWSN; middle: iterated M-OWNS ($\alpha_0 = \alpha$); bottom: fixed-carrier M-OWNS ($\alpha_0 = \omega$). Each at $n_x = 1000, 3000, 6000$ and 9000 .




Fig. 20 Rayleigh quotient α_r for the randomised inlet forcing, with Mode S and Mode F Rayleigh quotients from OWSN with LST initialisation. The fixed-carrier M-OWNS curve shows $\alpha_0 = \omega$ to indicate that no iterations are used. All randomised forcing calculations are converged at $n_x = 9000$.

Across all test cases in §5, the randomised-inlet study shows the smallest reduction in n_x at equivalent accuracy; the other cases show larger reductions. All computations reported here were performed on standard desktop hardware.

6 Conclusions

This paper has presented M-OWNS, a spatial marching method that combines the carrier-wave factoring of the parabolised stability equations (PSE) with the one-way projection of the recursive one-way Navier–Stokes (OWNS-R) framework. We have shown that this additional step offers significant computational benefits, while




Fig. 21 Pressure field $\text{Re}(p)$ for fixed-carrier M-OWNS ($\alpha_0 = \omega$) at $n_x = 3000$. White dashed line: boundary-layer edge.

still providing equivalent modelling fidelity to standard OWNS computations. The method is validated across subsonic and hypersonic boundary-layer flows as well as for three-dimensional crossflow disturbances. In addition, four forcing scenarios (at inlet or the wall) in a Mach 4.5 boundary-layer, provide results consistent with conventional OWNS and LHNS models.

The resolution advantage is explained by the spectral resolution comparison derived in §2.4: the carrier-wave factoring translates the integrator’s accuracy region from the origin to the carrier wavenumber α_0 in spectral space, reducing the residual wavenumbers of all modes near α_0 . At a single forcing frequency, the physically relevant eigenvalues cluster near the vorticity/entropy branch endpoint α_1 , so the criterion is generically satisfied. To leading order the spectral resolution condition is set by the eigenvalue geometry (§2.4); higher-order corrections depend on the scheme. In practice, wave-factoring allows the one-way projection to suppress the upstream-propagating modes at coarser streamwise step size Δx , a role the regularisation strategies of PSE cannot play without introducing characteristic errors [10].

For the fixed-carrier variant, M-OWNS reduces the total solve count by factors of two to eight in all test cases considered. The lower bound corresponds to the randomised-inlet case (§5.3.4), the most stringent multi-modal test: OWNS requires $n_x = 6000$ while fixed-carrier M-OWNS converges at $n_x = 3000$ with $N_\beta = 60$. The upper bound corresponds to the swept-cylinder configuration (§5.2), where fixed-carrier M-OWNS converges at $n_x = 1000$ against $n_x = 8000$ for OWNS. When a single mode dominates, the carrier-wave approach achieves even further reductions. For the swept-cylinder study, iterating M-OWNS converges with as few as 150 (n_x) grid points: a 47-fold reduction in computational cost relative to pure OWNS.

When the iterated Rayleigh quotient is used as a closure condition, a poorly chosen or divergent α_0 can degrade accuracy by displacing the accuracy region away from the excited spectrum (§2.5). The non-iterating variant, with α_0 set to the vorticity/entropy branch endpoint α_1 from Eq. (B.1), eliminates this risk at identical per-step cost to unfactored OWNS. The factoring provides no benefit when α_1 itself is near the origin; in this case $\alpha_0 = 0$ recovers unfactored OWNS with no penalty.

The choice $\alpha_0 = \alpha_1$ is appropriate when the excited spectrum is dominated by the vorticity/entropy branch and nearby discrete modes, which covers all test cases reported here. For regimes in which the acoustic and vortical responses are of comparable amplitude (*e.g.* optimal-disturbance and resolvent analyses in which both branches contribute to the gain) a single branch-endpoint carrier need not be optimal, and a minimax placement that balances the residual wavenumbers across the relevant branch endpoints may be preferable. More extensive investigations using the M-OWNS approach to modelling forced problems summarised in this paper may be found in Badcock [37].

Acknowledgements We thank the anonymous reviewers for the care and detail of their comments and suggestions. Their comments improved the rigour of the analysis and the overall clarity of the paper. We also thank Dr Joe Coppin (DSTL) for his continued support of this research.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funding This work was supported by the United Kingdom Defence Science and Technology Laboratory (DSTL), under contract DSTLX1000152974.

Declarations

Data Availability Statement The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Competing interests The authors declare no competing interests.

Appendix A Body-fitted compressible Navier–Stokes equations

We work in a body-fitted orthogonal coordinate system (x, y, z) , where x is the streamwise direction aligned normal to a wing's leading edge, y is wall-normal, and z is spanwise. The metric coefficient

$$\chi(x, y) = \frac{1}{1 - y \kappa(x)}, \quad (\text{A.1})$$

accounts for streamwise curvature $\kappa = \theta_x$, with θ the angle between the Cartesian x_1 axis and x . Spatial dimensions are normalised by the boundary-layer thickness at the initial streamwise station. Terms arising from the streamwise variation of the curvature, κ_x , are neglected.

The equations below are written in terms of the total (instantaneous) fields $(\rho, u, v, w, p, T, \mu)$, denoting density, velocity components, pressure, temperature, and dynamic viscosity. The corresponding state vector is

$$q^{\text{tot}} = (p, u, v, w, T)^{\top}, \quad (\text{A.2})$$

with ρ recovered from the equation of state.

For an ideal gas with ratio of specific heats γ , freestream Mach number M , Reynolds number R (based on the same reference length), Prandtl number σ , and bulk-to-shear viscosity ratio m , the non-dimensional equations read:

$$\rho_t + \chi u \rho_x + v \rho_y + w \rho_z + \rho(\chi u_x + v_y + w_z - \kappa \chi v) = 0, \quad (\text{A.3})$$

$$\begin{aligned} \rho(u_t + \chi u u_x + v u_y + w u_z - \kappa \chi u v) &= -\chi p_x \\ &+ \frac{1}{R} \left[\mu_x [(m+2)\chi^2(u_x - \kappa v) + m\chi(v_y + w_z)] \right. \\ &+ \mu_y [u_y + \chi(v_x + \kappa u)] + \mu_z (u_z + \chi w_x) \\ &+ \mu [(m+2)\chi^2(u_{xx} - \kappa v_x) + u_{yy} - \kappa \chi u_y + u_{zz} + (m+1)\chi(v_{xy} + w_{xz}) \\ &\left. - \kappa \chi^2(v_x + \kappa u)] \right], \end{aligned} \quad (\text{A.4})$$

$$\begin{aligned} \rho(v_t + \chi u v_x + v v_y + w v_z + \kappa \chi u^2) &= -p_y \\ &+ \frac{1}{R} \left[\mu_x [\chi u_y + \chi^2(v_x + \kappa u)] \right. \\ &+ \mu_y [(m+2)v_y + m\chi(u_x - \kappa v) + m w_z] + \mu_z (w_y + v_z) \\ &+ \mu [\chi^2 v_{xx} + (m+2)(v_{yy} - \kappa \chi v_y) + v_{zz} + (m+2)\kappa \chi^2(u_x - \kappa v) + \kappa \chi^2 u_x \\ &\left. + (m+1)(\chi u_{xy} + w_{yz}) \right], \end{aligned} \quad (\text{A.5})$$

$$\begin{aligned}
\rho (w_t + \chi u w_x + v w_y + w w_z) &= -p_z \\
&+ \frac{1}{R} \left[\mu_x (\chi u_z + \chi^2 w_x) + \mu_y (v_z + w_y) + \mu_z [(m+2)w_z + m v_y + m \chi (u_x - \kappa v)] \right. \\
&+ \mu [\chi^2 w_{xx} + w_{yy} + (m+2)w_{zz} - \kappa \chi w_y \\
&\quad \left. + (m+1)v_{yz} + (m+1)\chi(u_{xz} - \kappa v_z)] \right],
\end{aligned} \tag{A.6}$$

$$\begin{aligned}
\rho (T_t + \chi u T_x + v T_y + w T_z) &= (\gamma - 1)M^2 (p_t + \chi u p_x + v p_y + w p_z) \\
&+ \frac{1}{\sigma R} (\chi^2 \mu_x T_x + \mu_y T_y + \mu_z T_z + \mu (\chi^2 T_{xx} + T_{yy} + T_{zz} - \kappa \chi T_y)) \\
&+ \frac{(\gamma - 1)M^2}{R} \mu \mathcal{D},
\end{aligned} \tag{A.7}$$

where the viscous dissipation is

$$\begin{aligned}
\mathcal{D} &= (m+2) \left(w_z^2 + [v_y + \chi(u_x - \kappa v)]^2 \right) + 2m(v_y w_z + \chi(u_x w_z - \kappa v w_z)) \\
&+ 2(\chi u_z w_x + \chi u_y v_x + v_z w_y) + \chi^2 (v_x^2 + w_x^2) + 4\chi(\kappa v v_y - u_x v_y) \\
&+ 2\kappa \chi u(u_y + \chi v_x) + u_y^2 + w_y^2 + u_z^2 + v_z^2 + \kappa^2 \chi^2 u^2.
\end{aligned} \tag{A.8}$$

The system is closed by the ideal gas law

$$p = \frac{T\rho}{\gamma M^2}, \tag{A.9}$$

and Sutherland's law $\mu = \mu(T)$.

The linearised equations used in §2.1 follow from the decomposition $q^{\text{tot}} = Q + q$, where Q is the base flow and $q = (p', u', v', w', T')^\top$ is the perturbation. Primes are dropped in §2.1 onwards, so that $q = (p, u, v, w, T)^\top$ denotes the perturbation throughout.

Appendix B Recursion parameter selection

Recursion parameter selection for OWNS-R is more demanding than for other OWNS formulations. Sleeman and Colonius [14] showed that heuristic parameters yielding a stable OWNS-Projection [28] march can produce an unstable OWNS-R [13] march. Both retention and removal must therefore be controlled independently, making explicit parameter expressions essential for reproducibility. This appendix provides the complete formulae used for the subsonic and supersonic configurations in this work. Here, the spectrum of $\mathbf{M}$ at each streamwise location is computed via the QZ algorithm.

B.1 Spectral structure

The continuous spectral branches of Eq. (2) are obtained by frozen-coefficient analysis with freestream values. Substituting $\partial/\partial x \rightarrow i\alpha$ and $\partial/\partial y \rightarrow i\eta$ into Eq. (2) yields

$$\begin{aligned}
\alpha_1 &= \frac{\omega - \beta W_\infty}{U_\infty}, \quad \alpha_{2,3}(\eta) = \frac{\omega - \beta W_\infty}{U_\infty} + \frac{i(\beta^2 + \eta^2)}{U_\infty R}, \\
\alpha_{4,5}(\eta) &= \frac{M^2 U_\infty (\omega - \beta W_\infty) \pm \sqrt{\mu(\eta)}}{M^2 U_\infty^2 - T_\infty},
\end{aligned} \tag{B.1}$$

where

$$\mu(\eta) = T_\infty [(M^2 U_\infty^2 - T_\infty)(\beta^2 + \eta^2) + M^2 (\omega - \beta W_\infty)^2]. \tag{B.2}$$




Fig. 22 Numerical spectrum $\sigma(\mathbf{M})$ and parabolised spectrum $\sigma(\mathbf{P}_{N_\beta} \mathbf{M})$ for a 2D flat-plate boundary layer ($M = 0.02$, $R = 400$, $F = 86$, $b = 0$, $N_\beta = 30$, $c = 1$). Blue and red diamonds mark the recursion parameters β^+ and β^- ; green diamonds mark the roots β^* of the characteristic polynomial, lying near the $|\mathcal{P}_{N_\beta}| = 0.5$ contour. Teal regions ($|\mathcal{P}_{N_\beta}| \approx 1$) indicate retained modes; pink regions ($|\mathcal{P}_{N_\beta}| \approx 0$) indicate suppressed modes.

and U_∞ , W_∞ and T_∞ denote the freestream streamwise velocity, crossflow velocity and temperature respectively, each non-dimensionalised by the boundary-layer edge values at the inlet. The branches $\alpha_{1,2,3}$ represent vorticity/entropy modes; $\alpha_{4,5}$ are acoustic. In subsonic flows, α_{1-4} propagate downstream and α_5 propagates upstream; in supersonic flows all five propagate downstream [38, 39]. The critical spectral locations are

$$d_1 = \alpha_1, \quad d_2 = \frac{\alpha_4 + \alpha_5}{2}, \quad d_{3,4} = \alpha_{4,5}(0), \quad (\text{B.3})$$

where d_1 is the vorticity/entropy branch endpoint, d_2 is the acoustic branch midpoint, and $d_{3,4}$ are the propagating acoustic endpoints. The wall-normal grid stretching produces a non-uniform sampling of these branches: eigenvalues cluster near the branch endpoints (small $|\eta|$) and become sparse further along each branch (large $|\eta|$). The recursion parameters must follow this distribution.

B.2 Subsonic flows

The acoustic branches exhibit rotational symmetry about d_2 . We exploit this by placing β^+ to capture downstream modes and obtaining β^- through rotation by π about d_2 :

$$\beta_k^- = -(\beta_k^+ - d_2) + d_2. \quad (\text{B.4})$$

This ensures balanced filtering strength on both sides of the spectrum.

The N_β parameters are partitioned as

$$N_v = \lceil N_\beta/3 \rceil + 2, \quad N_a = N_\beta - N_v, \quad (\text{B.5})$$

allocating N_v to the vorticity/entropy branches and N_a to the acoustics, with denser placement near the real axis where the discrete eigenvalues concentrate. The downstream parameters are

$$\beta_k^+ = \begin{cases} d_1 + i|\alpha| y_{\text{Malik}}\left(\frac{k-1}{N_v-3}, 1, 0.4\right) & k = 1, \dots, N_v - 2, \\ \alpha & k = N_v - 1, \\ d_1 + 0.6(\alpha - d_1) & k = N_v, \\ d_2 + i\alpha_{\max} y_{\text{Malik}}\left(\frac{k-N_v-1}{N_a-1}, 1, 0.12\right) + p & k = N_v + 1, \dots, N_\beta, \end{cases} \quad (\text{B.6})$$




Fig. 23 Recursion parameter placement for subsonic 2D boundary-layers at $R_\delta = 400$, $b = 0$, $W_\infty = 0$, $F = 86$ with variable Mach number.

where α denotes the Rayleigh quotient (Eq. (9)) of the disturbance (or equally $|d_1|$ if in a strongly non-modal regime), y_{Malik} implements the non-uniform spacing of Eq. (41), and $\alpha_{\text{max}} = 28$ bounds the acoustic branch extent. The adjustment parameter

$$p = -5 \operatorname{sign}(d_1) \times 10^{-2} + 7.5 i \times 10^{-3} \quad (\text{B.7})$$

serves dual purposes: the imaginary component shifts β^- away from $\sigma(\mathbf{M})$ to prevent singularities in $\mathbf{P}_{N_\beta}$, while the real component clamps the contour $|\mathcal{P}_{N_\beta}| = 0.5$ so that it passes through the origin parallel to the real axis, separating the competing spectral regions.

The parameters are recalculated at each streamwise station as the base flow evolves and the spectrum migrates. The discrete mode (T-S or crossflow wave) must be tracked throughout the march; failure to do so leads to convergence difficulties or spurious migration to alternative eigenfunctions within $\sigma(\mathbf{M})$.

Incompressible two-dimensional flow Figure 22 illustrates the parameter placement for a flat-plate boundary-layer ($M = 0.02$, $R = 400$, $F = 86$, $b = 0$). This is the most challenging configuration: the downstream and upstream acoustic branches intersect near the origin, forcing β^+ and β^- parameters into direct competition. The rotational symmetry about d_2 resolves this by constructing fictitious upstream counterparts to the vorticity/entropy branches, balancing the filtering strength and preventing any single parameter from dominating.

Variable Mach number Figure 23 shows the effect of increasing Mach number at fixed $b = 0$. As M increases, the real parts of the acoustic branches $\alpha_{4,5}$ become more pronounced, progressively separating the downstream and upstream spectra and easing the parameter selection.

Variable spanwise wavenumber Figure 24 shows that increasing b separates the acoustic branches along the imaginary axis, eliminating the competition between β^+ and β^- parameters entirely. Parameter selection becomes straightforward for $b \gtrsim 0.1$.




Fig. 24 Recursion parameter placement for subsonic boundary-layers at $R_\delta = 400$, $M = 0.02$, $W_\infty = 0$, $F = 86$ with $b = 0.1$ and $b = 0.2$.




Fig. 25 Spectral evolution for the swept cylinder ($R = 132$, $\beta = 0.4$, $N_\beta = 20$). As R_δ increases, W_∞/U_∞ decreases from 3.79 to 0.960 and the vorticity/entropy branches migrate toward the acoustic spectrum. The same subsonic parameter formulas (Eqs. B.6–B.4) handle this evolution without modification.

Swept cylinder The swept cylinder flow ($R = 132$, $\beta = 0.4$) introduces an evolving spectral structure absent in flat-plate cases: the crossflow ratio W_∞/U_∞ decreases downstream, causing the vorticity/entropy branches to migrate toward the acoustic spectrum (Fig. 25). Equations (B.6) and (B.4) accommodate this through two mechanisms. The symmetric placement of $\beta^\pm$ about d_2 maintains balanced filtering despite the migration, and the $|\alpha|$ scaling factor in Eq. (B.6) contracts the vorticity/entropy parameter extent as the branches shorten downstream. The nonzero spanwise wavenumber keeps the acoustic branches well separated throughout, so the parameter selection remains robust despite the spectral evolution.

These parameters have been validated across wall-normal discretisations (n_y from 71 to 1201) and with the tuning parameter $c \in [0.2, 10]$.

B.3 Hypersonic flows

In supersonic and hypersonic boundary-layers all continuous branches propagate downstream. Upstream propagation occurs only through discrete acoustic modes, trapped between the wall and sonic line, that populate the negative imaginary half-plane with increasing density [39]. The rotational symmetry exploited above no longer applies, and the β^+ and β^- parameters require independent placement strategies.




Fig. 26 Numerical spectrum $\sigma(\mathbf{M})$ and parabolised spectrum $\sigma(\mathbf{P}_{N_\beta}\mathbf{M})$ for a supersonic flat-plate boundary-layer ($M = 4.5$, $R_\delta = 400$, $N_\beta = 30$, $c = 0.5$).

The downstream parameters are positioned along the vorticity/entropy branches:

$$\beta_j^+ = d_1 + i y_{\text{Malik}} \left(\frac{j-1}{N_\beta-1}, 5, 0.2 \right). \quad (\text{B.8})$$

The upstream parameters target the discrete modes through enhanced clustering near the real axis:

$$\beta_j^- = -0.15 - 0.01i - \left(10(-1)^j + 50i \right) y_{\text{Malik}} \left(\frac{j-1}{N_\beta-1}, 1, 0.01 \right). \quad (\text{B.9})$$

The alternating term $(-1)^j$ provides lateral spreading while maintaining dense coverage near $\text{Im}(\alpha) = 0$. A β^+ can additionally be placed on the Rayleigh quotient if the quantity is well-defined and inside the convex hull of the dominant spectrum.

Figure 26 confirms that this configuration ($N_\beta = 30$, $c = 0.5$, optimised for $M = 4.5$) retains the complete downstream spectrum while eliminating upstream propagation. Modified versions of these parameters have been applied to transonic, supersonic ($M = 1.6$) and an additional hypersonic ($M = 6$) case.

References

1. Mack, L.M.: Boundary-layer linear stability theory. In: AGARD Special Course on Stability and Transition of Laminar Flow, pp. 3–81. AGARD Report R-709, Neuilly-sur-Seine, France (1984)
2. Theofilis, V.: Global linear instability. *Annu. Rev. Fluid Mech.* **43**, 319–352 (2011). <https://doi.org/10.1146/annurev-fluid-122109-160705>
3. Appel, T.: Boundary layer instabilities due to surface irregularities: A harmonic Navier-Stokes approach, Imperial College London (2021). <https://doi.org/10.25560/88722>. (PhD thesis)
4. Herbert, T.: Parabolized Stability Equations. *Annu. Rev. Fluid Mech.* **29**(1), 245–283 (1997). <https://doi.org/10.1146/annurev.fluid.29.1.245>
5. Bertolotti, F.P., Herbert, T.: Analysis of the linear stability of compressible boundary layers using the PSE. *Theoret. Comput. Fluid Dyn.* **3**(2), 117–124 (1991). <https://doi.org/10.1007/BF00271620>
6. Turner, M.R., Hammerton, P.W.: Asymptotic receptivity analysis and the parabolized stability equation: a combined approach to boundary layer transition. *J. Fluid Mech.* **562**, 355–381 (2006). <https://doi.org/10.1017/S0022112006001108>
7. Li, F., Malik, M.R.: On the nature of PSE approximation. *Theoret. Comput. Fluid Dyn.* **8**, 253–273 (1996). <https://doi.org/10.1007/BF00639695>
8. Li, F., Malik, M.R.: Spectral analysis of parabolized stability equations. *Comput. Fluids* **26**(3), 279–297 (1997). [https://doi.org/10.1016/S0045-7930\(96\)00044-8](https://doi.org/10.1016/S0045-7930(96)00044-8)
9. Andersson, P., Henningson, D.S., Hanifi, A.: On a stabilization procedure for the parabolic stability equations. *J. Eng. Math.* **33**, 311–332 (1998). <https://doi.org/10.1023/A:1004367704897>
10. Towne, A., Rigas, G., Colonius, T.: A critical assessment of the parabolized stability equations. *Theoret. Comput. Fluid Dyn.* **33**, 359–382 (2019). <https://doi.org/10.1007/s00162-019-00498-8>
11. Towne, A., Colonius, T.: One-way spatial integration of hyperbolic equations. *J. Comput. Phys.* **300**, 844–861 (2015). <https://doi.org/10.1016/j.jcp.2015.08.015>

12. Rigas, G., Schmidt, O.T., Colonius, T., Brés, G.A.: One-way Navier–Stokes and resolvent analysis for modeling coherent structures in a supersonic turbulent jet. In: 23rd AIAA/CEAS Aeroacoustics Conference, Denver, CO, USA (2017). <https://doi.org/10.2514/6.2017-4046>. AIAA Paper 2017-4046
13. Zhu, M., Towne, A.: Recursive one-way Navier–Stokes equations with PSE-like cost. *J. Comput. Phys.* **473**, 111744 (2023). <https://doi.org/10.1016/j.jcp.2022.111744>
14. Sleeman, M.K., Colonius, T.: Greedy recursion parameter selection for the One-Way Navier–Stokes (OWNS) equations (2025). <https://arxiv.org/abs/2506.02320>
15. Badcock, E.J., Mughal, S.: Application of the One-Way Navier–Stokes equations to model disturbance evolution in laminar separation bubbles. In: Kato, K., Inasawa, A., Matsubara, M. (eds.) *Proceedings of the 10th IUTAM Symposium on Laminar–Turbulent Transition*, pp. 323–329. Springer, Singapore (2026). https://doi.org/10.1007/978-981-96-9829-5_42
16. Malik, M.R., Li, F., Chang, C.-L.: Crossflow disturbances in three-dimensional boundary layers: nonlinear development, wave interaction and secondary instability. *J. Fluid Mech.* **268**, 1–36 (1994). <https://doi.org/10.1017/S0022112094001242>
17. Mughal, M.S.: Active control of wave instabilities in three-dimensional compressible flows. *Theoret. Comput. Fluid Dyn.* **12**, 195–217 (1998). <https://doi.org/10.1007/s001620050106>
18. Haj-Hariri, H.: Characteristics analysis of the parabolized stability equations. *Stud. Appl. Math.* **92**(1), 41–53 (1994). <https://doi.org/10.1002/sapm199492141>
19. Stewart, G.W., Sun, J.-G.: *Matrix Perturbation Theory*, Academic Press, Boston (1990)
20. Chang, C.-L., Malik, M.R., Erlebacher, G., Hussaini, M.Y.: *Linear and nonlinear PSE for compressible boundary layers*, Hampton, VA (1993). (NASA-CR-191537 Technical Report)
21. Chu, B.-T.: On the energy transfer to small disturbances in fluid flow (Part I). *Acta Mech.* **1**(3), 215–234 (1965). <https://doi.org/10.1007/BF01387235>
22. Hanifi, A., Schmid, P.J., Henningson, D.S.: Transient growth in compressible boundary layer flow. *Phys. Fluids* **8**(3), 826–837 (1996). <https://doi.org/10.1063/1.868864>
23. Guo, Y., Malik, M., Chang, C.-L.: A solution adaptive approach for computation of linear waves. In: 13th AIAA Computational Fluid Dynamics Conference, Snowmass Village, CO, USA (1997). <https://doi.org/10.2514/6.1997-2072>. AIAA Paper 1997-2072
24. Franco Sumariva, J.A., Hein, S.J.: Adaptive harmonic linearized Navier–Stokes equations used for boundary layer instability analysis in the presence of large streamwise gradients. In: AIAA Aerospace Sciences Meeting, Kissimmee, FL, USA (2018). <https://doi.org/10.2514/6.2018-1548>. AIAA Paper 2018-1548
25. Herbert, T.: Boundary-layer transition — analysis and prediction revisited. In: 29th Aerospace Sciences Meeting. American Institute of Aeronautics and Astronautics, Reno, NV, USA (1991). <https://doi.org/10.2514/6.1991-737>. AIAA Paper 1991-737
26. Briggs, R.J.: *Electron-Stream Interaction with Plasmas*. The MIT Press, Cambridge, MA (1964). <https://doi.org/10.7551/mitpress/2675.001.0001>
27. Towne, A.S.: Advancements in jet turbulence and noise modeling: Accurate one-way solutions and empirical evaluation of the nonlinear forcing of wavepackets. PhD thesis, California Institute of Technology (2016). <https://doi.org/10.7907/Z99884XJ>
28. Towne, A., Rigas, G., Kamal, O., Pickering, E., Colonius, T.: Efficient global resolvent analysis via the one-way Navier–Stokes equations. *J. Fluid Mech.* **948** (2022). <https://doi.org/10.1017/jfm.2022.647>
29. Rudel, C., Pernet, S., Brazier, J.-P.: Numerical factorization of propagation operator for hyperbolic equations and application to one-way, true amplitude one-way equations and Bremmer series. *J. Sci. Comput.* **93**(1), 27 (2022). <https://doi.org/10.1007/s10915-022-01985-7>
30. Malik, M.R.: Numerical methods for hypersonic boundary layer stability. *J. Comput. Phys.* **86**(2), 376–413 (1990). [https://doi.org/10.1016/0021-9991\(90\)90106-B](https://doi.org/10.1016/0021-9991(90)90106-B)
31. Thompson, K.W.: Time dependent boundary conditions for hyperbolic systems. *J. Comput. Phys.* **68**(1), 1–24 (1987). [https://doi.org/10.1016/0021-9991\(87\)90041-6](https://doi.org/10.1016/0021-9991(87)90041-6)
32. Fedorov, A.V.: Receptivity of a high-speed boundary layer to acoustic disturbances. *J. Fluid Mech.* **491**, 101–129 (2003). <https://doi.org/10.1017/S0022112003005263>
33. Fedorov, A., Tumin, A.: Branching of discrete modes in high-speed boundary layers and terminology issues. In: 40th Fluid Dynamics Conference and Exhibit. AIAA Paper 2010-5003, Chicago, Illinois (2010). <https://doi.org/10.2514/6.2010-5003>
34. Ma, Y., Zhong, X.: Receptivity of a supersonic boundary layer over a flat plate .Part 1. Wave structures and interactions. *J. Fluid Mech.* **488**, 31–78 (2003). <https://doi.org/10.1017/S0022112003004786>
35. Ma, Y., Zhong, X.: Numerical simulation of receptivity and stability of nonequilibrium reacting hypersonic boundary layers. In: 39th Aerospace Sciences Meeting and Exhibit. AIAA Paper 2001-892, Reno, NV, USA (2001). <https://doi.org/10.2514/6.2001-892>
36. Ma, Y., Zhong, X.: Receptivity of a supersonic boundary layer over a flat plate. Part 2. Receptivity to free-stream sound. *J. Fluid Mech.* **488**, 79–121 (2003). <https://doi.org/10.1017/S0022112003004798>
37. Badcock, E.J.: *Modal One-Way Navier–Stokes: Efficient spatial marching for modal and non-modal boundary-layer disturbances*. In: Embargoed; available via institutional repository, Imperial College London (2026). (PhD thesis)
38. Balakumar, P., Malik, M.R.: Discrete modes and continuous spectra in supersonic boundary layers. *J. Fluid Mech.* **239**, 631–656 (1992). <https://doi.org/10.1017/S0022112092004555>
39. Fedorov, A.: Transition and stability of high-speed boundary layers. *Annu. Rev. Fluid Mech.* **43**(1), 79–95 (2011). <https://doi.org/10.1146/annurev-fluid-122109-160750>