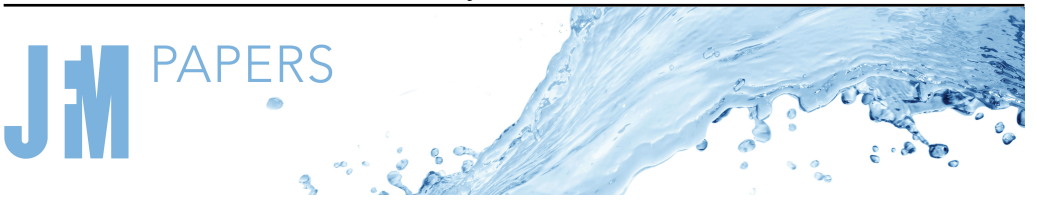
J. Fluid Mech. (2026), *vol.* 1036, A63, doi:10.1017/jfm.2026.11671



Data-driven nonlinear aerodynamics models with certifiably optimal boundedness properties

A. Leonid Heide[1], **Shih-Chi Liao**[2], **Sergio Castiblanco-Ballesteros**[3], **Gustaaf Jacobs**[3], **Peter Seiler**[2] **and Maziar S. Hemati**[1]

[1]Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN, USA

[2]Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, USA

[3]Department of Aerospace Engineering, San Diego State University, San Diego, CA, USA

Corresponding author: A. Leonid Heide, heide116@umn.edu

(Received 22 August 2025 UTC; revised 1 April 2026 UTC; accepted 30 April 2026 UTC)

Obtaining predictive low-order models is a central challenge in fluid dynamics. Data-driven frameworks have been widely used to obtain low-order models of aerodynamic systems; yet, resulting models tend to yield predictions that grow unbounded with time. Recently introduced stability-promoting methods can facilitate the identification of bounded models, but tend to require extensive brute-force tuning even in the context of simple academic systems. Here, we show how recent theoretical advances in the long-term boundedness of dynamical systems can be integrated into data-driven modelling frameworks to ensure that resulting models will yield bounded predictions of incompressible flows. Specifically, we propose to solve a particular set of convex semidefinite programming problems to (i) certify whether a system admits a globally attracting bounded set for the chosen modelling parameters, and (ii) compute a model with the optimal (tightest) bound on this globally attracting set. We demonstrate the approach via integration within the sparse identification of nonlinear dynamics (SINDy) modelling framework. Application on two low-order benchmark problems establishes the merits of the approach. We then apply our approach to obtain a low-order (six-mode) model of unsteady separation over a NACA-65(1)-412 aerofoil at $Re = 20\,000$ – a flow that has been notoriously difficult to model using data-driven methods. The resulting model is found to accurately predict the dynamics of unsteady separation, with model predictions remaining bounded indefinitely. We anticipate this work will benefit future efforts in modelling strongly nonlinear flows, especially in settings where physically viable long-term forecasts are paramount.

Key words: low-dimensional models, computational methods, control theory

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

1. Introduction

Low-order models are an enabling tool in fluid dynamics, providing unique opportunities for rapid prediction and forecasting, estimation and control, and design optimisation (Benner, Gugercin & Willcox 2015; Brunton & Noack 2015; Rowley & Dawson 2017). However, most use-cases require that low-order models remain faithful to the physics and underlying dynamics of the flows under consideration. This can be especially challenging in aerodynamics applications because the difficulty of crafting a low-order model that faithfully captures the essential physical and dynamical processes of a flow grows with increasing Reynolds number (Moehlis, Faisst & Eckhardt 2004; Cavalieri & Nogueira 2022).

In fluid dynamics, low-order models have traditionally been obtained within the framework of projection-based model reduction (Holmes *et al.* 2012), whereby the Navier–Stokes equations are projected onto a low-dimensional subspace (e.g. Galerkin projection onto a subspace spanned by dominant proper orthogonal decomposition (POD) modes). Indeed, projection-based reduced-order models (ROMs) have proven a powerful tool for distilling the complex dynamics of unsteady flows into a compact set of ordinary differential equations (ODEs) that support real-time analysis, optimisation and predictive control (Rowley & Dawson 2017; Taira *et al.* 2017; Brunton, Noack & Koumoutsakos 2020; Taira *et al.* 2020).

Despite demonstrated successes in obtaining simplified yet physically faithful models in fluid dynamics, projection-based model reduction can be difficult to implement in practice. Projection-based approaches often demand full access to the governing equations and their discretisations – an intrusive requirement that limits their applicability, especially when working with complex computational fluid dynamics (CFD) codes (Peherstorfer & Willcox 2016). Projection-based model reduction can also result in models that diverge due to neglected energy transfers (Cazemier, Verstappen & Veldman 1998; Schlegel & Noack 2015).

In contrast, data-driven frameworks – such as sparse identification of nonlinear dynamics (SINDy) (Brunton *et al.* 2016; Loiseau & Brunton 2018), physics-informed neural networks (Raissi, Perdikaris & Karniadakis 2019; Thuerey *et al.* 2020), operator learning (Lu *et al.* 2021) and operator inference (Peherstorfer & Willcox 2016; Kramer, Peherstorfer & Willcox 2024) – are non-intrusive, inferring compact dynamical systems directly from data generated by CFD or measured in experiments. The data-driven approach enables rapid model construction and deployment even when a first-principles description is unavailable or only partially observed (Brunton *et al.* 2016; Raissi *et al.* 2019).

Nonetheless, low-order models obtained from data-driven frameworks have their own limitations as well, often failing to generalise outside training datasets and commonly yielding unphysical energy growth under extrapolation (Duraisamy, Iaccarino & Xiao 2019; Karniadakis *et al.* 2021). As such, much work has focused on ensuring that data-driven models adhere to known physical laws (e.g. conservation of mass, momentum and energy). In the case of incompressible flows, the quadratic nature of the convective nonlinearity in the Navier–Stokes equations admits energy transfer between modal triads – so-called triadic interactions (Schmidt 2020). Furthermore, the divergence-free condition on the velocity field along with mild assumptions regarding boundary conditions results in the convective nonlinearity being lossless – i.e. the quadratic terms are energy preserving, only serving to redistribute kinetic energy without net production or dissipation. Data-driven methods have been developed to impose these physics-based constraints within modelling frameworks (Loiseau & Brunton 2018); however, the enforcement of lossless

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

quadratic terms is not sufficient to guarantee that resulting models will maintain physically viable predictions over long-time horizons. Indeed, it is common for such models to yield predictions that grow unbounded with time – an indicator of unphysical behaviour. This is particularly problematic in flow control and estimation applications: a ROM that is not globally bounded may diverge or exit its regime of validity, undermining any controller or estimator based on it.

Notions of stability and boundedness have been investigated and developed to contend with the possibility of this unphysical feature in the dynamics of low-order models. Methods have been proposed for efficiently analysing the asymptotic stability of quadratic ROMs (Kramer 2021; Enayati Kafshgarkolaei & Hemati 2025) or for directly identifying asymptotically stable ROMs from data (Goyal, Duff & Benner 2024). However, asymptotic stability is often a stronger condition than would be desired in many aerodynamics applications where trajectories remain bounded with time, but do not necessarily asymptote to an equilibrium point (e.g. wake shedding from an aerofoil). For such systems, it is natural to consider notions of long-time boundedness of trajectories (Khalil 2002).

Prior work by Schlegel & Noack (2015) studied the problem of long-time boundedness and established that quadratic dynamical systems whose quadratic nonlinearity is lossless can admit a trapping region – a monotonically attracting invariant set within which all trajectories remain indefinitely. The Schlegel and Noack trapping theorem provides conditions for the existence of a trapping region, which can be used to prove the long-time boundedness of projection-based ROMs for incompressible flows via non-convex optimisation (Schlegel & Noack 2015). The theorem also provides a (conservative) upper bound on the size of the trapping region, if one exists. Recent work by Kaptanoglu *et al.* (2021) has leveraged the Schlegel and Noack trapping theorem within the SINDy modelling framework to facilitate the identification of bounded models from data. This trapping SINDy approach resorts to using gradient descent and relax-and-split methods (Champion *et al.* 2020) for the numerical solution of the modelling problem. The authors have provided a Python implementation of their method, which includes several benchmark examples that demonstrate the method's efficacy in yielding bounded models from data (de Silva *et al.* 2020; Kaptanoglu *et al.* 2022). Recent work by Peng *et al.* (2024) have extended the trapping SINDy framework to consider quadratic systems whose nonlinearity weakly violates the lossless constraint, using local notions of long-time boundedness to enable modelling of a larger class of systems.

Alongside these ROM-focused developments, sum-of-squares (SOS) and semidefinite programming (SDP) have also been used in fluid mechanics to certify stability and design feedback control laws for polynomial (or polynomially approximated) flow models (Goulart & Chernyshenko 2012; Chernyshenko *et al.* 2014; Huang *et al.* 2015; Lasagna *et al.* 2016). In related dynamical-systems literature, polynomial optimisation and SDP relaxations have been developed to bound extrema over global attractors and to compute converging outer approximations of attractors for polynomial systems (Goluskin 2020; Schlosser & Korda 2021). More generally, convex optimisation has been advocated as a systematic way to learn polynomial dynamical systems subject to side information, with Lyapunov-type inequalities (Ahmadi & El Khadir 2023). The trapping-region viewpoint adopted in the present work can be interpreted as a particular Lyapunov-based boundedness certificate within this broader framework, closely related to classical energy estimates for fluid models (Doering & Gibbon 1995; Khalil 2002).

In our recent work (Liao *et al.* 2025), we have made theoretical advances that improve upon the Schlegel and Noack trapping theorem and establish a pair of convex SDP conditions that have been shown to yield trapping regions that are significantly smaller

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

than those from earlier analytical estimates. The first SDP poses a feasibility problem to certify whether a monotonically attracting trapping region exists. If this problem is infeasible, one can formally conclude that no trapping region is admissible for the given system.

When the feasibility condition is satisfied, a second (dual) SDP is solved to minimise the radius of the invariant ball, thereby proving that the dynamical system is confined by the smallest possible Euclidean sphere. Since there is a unique optimal solution for the radius of the smallest invariant ball – by virtue of convexity – it follows that this optimal trapping-region size can provide a distinct dynamical boundedness property that should be captured by a low-order model in order to be dynamically consistent with the governing equations, in as much as any physical property should be captured for a model to be physically consistent with the governing equations.

In this paper we show how the trapping-SDP conditions from (Liao *et al.* 2025) can be integrated within data-driven modelling frameworks to formally guarantee long-time boundedness in the modelling process, and to (at least approximately) achieve dynamical consistency as needed to produce physically viable predictions over long-time horizons. In the context of prior SOS/SDP approaches, the goal of the present paper is not to develop the most general polynomial Lyapunov analysis, but rather to integrate a pair of tractable convex certificates for Euclidean trapping balls into a data-driven identification loop tailored to lossless quadratic ROMs. To this end, our proposed trapping-SDP modelling framework consists of an alternating (block coordinate descent) minimisation that cycles between the solution of two convex optimisation problems – (i) a regression problem, and (ii) a stability-enforcement problem – both of which can be solved efficiently using available convex optimisation software packages.

To focus the discussion, we demonstrate this integration within the SINDy modelling framework – though integration within alternative data-driven modelling frameworks follows similarly. The proposed trapping-SDP modelling approach is first presented on two canonical examples that allow direct comparison with known governing equations and associated optimal trapping regions. We then apply our approach to obtain a six-mode ROM of unsteady flow separation over a NACA 65(1)-412 aerofoil at a Reynolds number of $Re = 20\,000$ and angle of attack $\alpha = 4^\circ$ based on data from direct numerical simulations (DNS). In each scenario, we obtain models with guaranteed long-time boundedness properties and report the optimal trapping regions for the given set of parameters. All identified models are found to be predictive for short- and long-time horizons when applied to off-training validation tests.

We note that relative to trapping SINDy (Kaptanoglu *et al.* 2021), which couples the Schlegel–Noack trapping conditions to SINDy using relax-and-split and gradient-based iterations to handle a non-convex stabilisation step, we formulate boundedness enforcement as convex SDP subproblems that admit global solutions. Specifically, the feasibility SDP provides either (i) a certificate that a globally monotonically attracting Euclidean trapping ball exists for the chosen library and hyperparameters, or (ii) a rigorous infeasibility certificate within the adopted quadratic-energy certificate class. When feasible, we compute the tight trapping-ball radius (for a fixed centre) via a SDP dual, avoiding reliance on conservative analytic radius estimates based on worst-case eigenvalue bounds. Moreover, boundedness enters our identification procedure as a hard semidefinite constraint, so every accepted candidate model is accompanied by an explicit certificate, enabling systematic sweeps over hyperparameters rather than *ad hoc* tuning.

The paper is organised as follows. In § 2 we review requisite background on ROMs for incompressible flows, long-time boundedness, trapping regions, trapping SDPs and SINDy. In § 3 we present our trapping-SDP modelling framework within the context of

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

SINDy. Section 4 reports on the application of the trapping-SDP modelling framework on two numerical benchmark systems. Application to data-driven model reduction of separated flow is reported in § 5. Conclusions are summarised in § 6.

2. Background and preliminaries

2.1. *Reduced-order models of incompressible flows*

Before presenting our method, we first establish the essential system properties of ROMs for incompressible flows. Reduced-order modelling aims to approximate the (infinite-dimensional) partial differential equations (PDEs) that govern the system with a low-(finite-)dimensional representation. A common model reduction technique is the Galerkin projection, i.e. an orthogonal projection of these PDEs onto a finite-dimensional subspace spanned by a set of basis functions. In incompressible flows, the nonlinear (convective) term of the Navier–Stokes equations merely redistributes energy among the flow's degrees of freedom – a so-called lossless property of the convective nonlinearity. Galerkin projection of the incompressible Navier–Stokes equations onto an n-dimensional subspace yields ODEs of the form

$$\dot{x}(t) = c + Lx(t) + f(x(t)), \tag{2.1}$$

where $x(t) \in \mathbb{R}^n$ is the reduced state, the function, $f : \mathbb{R}^n \to \mathbb{R}^n$ is quadratic, with its ith entry given by $f_i(x) = x^\top Q^{(i)} x$ $(i = 1, \ldots, n)$, the matrix $L \in \mathbb{R}^{n\times n}$ contains the linear terms and $c \in \mathbb{R}^n$ is a constant vector. The quadratic term matrices $\{Q^{(i)}\}_{i=1}^n \in \mathbb{R}^{n\times n}$ here are symmetric. If the kinetic energy of the system (2.1) is defined as

$$K_0(x) = \frac{1}{2}\|x\|^2, \tag{2.2}$$

then the nonlinearity $f(x(t))$ is said to be lossless if

$$x^\top f(x) = 0 \quad \text{for all } x \in \mathbb{R}^n. \tag{2.3}$$

This indicates that f neither generates nor dissipates energy, but instead redistributes it among the components of x. Ensuring energy conservation requires that the quadratic term matrices $Q^{(i)}$ comprising the nonlinear term $f(x(t))$ satisfy the condition

$$\sum_{i,j,k=1}^{n} Q^{(i)}_{jk} x_i x_j x_k = 0 \quad \text{for all } x \in \mathbb{R}^n. \tag{2.4}$$

Under the lossless condition, the time derivative of the energy (2.2) along trajectories of the system (2.1) simplifies to

$$\frac{d}{dt} K_0(x(t)) = x(t)^\top \left(c + L_S x(t)\right), \tag{2.5}$$

where $L_S = (1/2)(L + L^\top)$ denotes the symmetric part of L. Equation (2.5) indicates that the quadratic nonlinearity f does not affect the energy evolution.

If the system's coordinates are translated by a constant vector $m \in \mathbb{R}^n$, where $y = x - m$, the system dynamics (2.1) become

$$\dot{y}(t) = d(m) + A(m)y(t) + f(y(t)), \tag{2.6}$$

where $d(m) = c + Lm + f(m)$ and the matrix $A(m)$ is given by

$$A(m) = L + 2\begin{pmatrix} m^\top Q^{(1)} \\ \vdots \\ m^\top Q^{(n)} \end{pmatrix}. \tag{2.7}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

The quadratic nonlinearity $f(y)$ in the shifted system (2.6) remains identical to that in the original system (2.1), ensuring that the lossless property is preserved after translation. The symmetric part of $A(m)$, denoted $A_S(m)$, is defined as

$$A_S(m) = \frac{1}{2}\big(A(m) + A(m)^\top\big) = L_S - \sum_{k=1}^{n} m_k Q^{(k)}, \tag{2.8}$$

where m_k is the kth element of the translation vector m. A more detailed derivation of (2.8) is provided in Appendix A

The energy in the shifted coordinates is $K_m(y) = (1/2)\|y\|^2$ and its time derivative is

$$\frac{d}{\mathrm{d}t} K_m(y(t)) = y(t)^\top d(m) + y(t)^\top A_s(m) y(t). \tag{2.9}$$

Further details on Galerkin projection and the derivation of lossless ROMs can be found in Appendix B and Schlegel & Noack (2015), Loiseau & Brunton (2018).

2.2. *Long-time boundedness and trapping regions*

In reduced-order modelling of fluid dynamics, ensuring solutions remain bounded is a crucial problem, as unbounded behaviours correspond to non-physical or unrealistic phenomena. Here, we focus on globally uniformly ultimately bounded systems (Khalil 2002), which means there exists a region such that trajectories from any initial condition eventually lie within a ball of radius $\beta > 0$. Formally, for a nonlinear system (2.1), the following holds for some $\beta > 0$:

$$\|x(t)\| \leqslant \beta \quad \text{for all } t \geqslant t_0 + T\big(\|x_0\|\big). \tag{2.10}$$

Here $T(\cdot)$ is a function depending on the initial condition x_0 but not on the initial time t_0.

A practical way to verify boundedness in lossless quadratic systems is through the concept of a trapping region (Schlegel & Noack 2015), a compact, forward-invariant set $D \subseteq \mathbb{R}^n$ that, once entered by a trajectory, cannot be exited. If the energy is strictly decreasing outside D (i.e. $\mathrm{d}K/\mathrm{d}t < 0$), the trapping region is said to be globally monotonically attracting, and hence, all trajectories eventually enter D. Here, D may be any compact forward-invariant set; it need not be spherical. In the remainder of this work we specialise to Euclidean balls because our boundedness certificates are constructed using the quadratic kinetic energy storage function, whose sublevel sets are precisely ℓ_2 balls. This choice yields a tractable certificate class and an interpretable bound, though it can be conservative relative to higher-degree SOS Lyapunov functions and generalised invariant sets (Goulart & Chernyshenko 2012; Chernyshenko *et al.* 2014; Huang *et al.* 2015; Lasagna *et al.* 2016; Goluskin 2020; Schlosser & Korda 2021). More general ellipsoidal bounds can be obtained by replacing $K_m(y)$ with a weighted quadratic storage function $K_{m,P}(y) = (1/2)y^\top P y$ with $P \succ 0$, whose sublevel sets are ellipsoids; however, we focus on Euclidean balls for simpler repeated embedding of the certificates inside a data-driven identification loop (Boyd *et al.* 1994).

Lorenz (1963) provided a condition for the existence of such a trapping region for lossless systems like (2.1), stating that if the symmetric part of the linear operator satisfies $L_S = (1/2)(L + L^\top) \prec 0$ then a trapping region exists. To see this, rewrite the energy of the system (2.5) in terms of the operator L_S as

$$\frac{d}{\mathrm{d}t} K_0\big(x(t)\big) = x(t)^\top c + x(t)^\top L_S\, x(t). \tag{2.11}$$

Then, it can be shown that if L_s is negative definite, there exists a finite radius outside of which the energy decreases, thus implying system trajectories will eventually enter a

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

finite ball from which they cannot exit – i.e. trajectories will remain bounded indefinitely. For a detailed derivation, we refer the reader to Lorenz (1963), Schlegel & Noack (2015). We note that these trapping-region arguments are naturally interpreted as quadratic Lyapunov (energy) estimates for dissipative systems, consistent with classical treatments of boundedness and absorbing sets (Doering & Gibbon 1995; Khalil 2002).

Schlegel & Noack (2015) further generalised this by leveraging the fact that quadratic systems preserve their nonlinearities under a coordinate shift, as described above in the introduction to § 2. Schlegel and Noack showed that if there is a region such that the shifted operator $A_s(m)$ is negative definite, a less conservative estimate for a globally monotonically attracting trapping region in the form of a ball $B\big(m,\, R_m\big)$ centred at m with radius R_m can be obtained. Moreover, Schlegel and Noack prove that the system (2.6) possesses a trapping region if and only if there exists an m such that

$$A_s(m) \;=\; L_s \;-\; \sum_{i=1}^{n} m_i \, Q^{(i)} \;\prec\; 0. \tag{2.12}$$

In other words, by choosing m such that $A_s(m)$ is negative definite, we ensure that the energy (2.9) decreases outside the trapping region. When (2.12) holds, the radius of the trapping region centred at m is given by

$$R_m \;=\; \frac{\|\, c + L\, m + f(m)\|}{\big|\lambda_{max}\big(A_s(m)\big)\big|}, \tag{2.13}$$

where $\lambda_{max}\big(A_s(m)\big)$ is the largest (least negative) eigenvalue of $A_s(m)$. The numerator $\|c + Lm + f(m)\|$ here essentially acts as a constant forcing or 'drift' term, while the denominator $|\lambda_{max}\big(A_s(m)\big)|$ is the least negative (i.e. worst-case) rate at which the quadratic term can remove energy from the system. All trajectories of the system eventually enter $B(m,\, R_m)$ and remain there, establishing global boundedness. As will be discussed in the next section, the radius R_m is an upper-bound estimate on the trapping-region size and is therefore usually overly conservative. For more detailed proofs and broader discussions of trapping regions in lossless quadratic systems, we refer the reader to Schlegel & Noack (2015), Liao *et al.* (2025).

2.3. *Optimal trapping regions via convex optimisation*

While the Schlegel-Noack condition shown in (2.12) gives a clear algebraic criterion for global boundedness, its practical implementation poses two challenges. First, verifying the trapping-region condition requires a search over the translation vector $m \in \mathbb{R}^n$ to render the symmetric part $A_s(m)$ negative definite. Schlegel & Noack (2015) cast this as a non-convex inf-sup problem, minimising the maximal eigenvalue of $A_s(m)$, and advocate simulated annealing for its solution. However, without convexity there is no guarantee of finding the true global minimum; failure of the annealing search therefore leaves the boundedness test inconclusive, since a valid trapping region may still exist.

Second, when a candidate shift vector m is found using the Schlegel and Noack method, the resulting radius shown in (2.13) is an upper-bound approximation, and thus, an inherently conservative overestimate. For a detailed explanation, we refer the reader to Liao *et al.* (2025). In short, this conservatism arises due to radius R_m being estimated from the worst-case alignment of the linear 'drift' vector $c + Lm + f(m)$ with the maximum (most conservative) eigenvalue of the quadratic 'dissipation' rate $\lambda_{max}(A_S(m))$. Consequently, the resulting bounds may be excessively large, which can be particularly problematic in reduced-order fluid dynamics models where unphysically large trapping regions do not accurately reflect the true system dynamics.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

In our recent work (Liao *et al.* 2025), we developed a fully convex framework to remove these limitations and both certify the existence of trapping regions and compute their tightest possible radii for lossless quadratic systems. We provide a summary of these optimisation problems here, and refer the reader to Liao *et al.* (2025) for detailed proofs and derivations.

To determine whether a trapping region exists for a given system, we recast the Schlegel–Noack criterion shown in (2.12) as a single linear matrix inequality (LMI). This is done by noting that in the condition (2.12), the shifted term $A_s(m)$ depends affinely on m. Checking negative definiteness of $A_s(m)$ is thus an LMI feasibility problem (Boyd & Vandenberghe 2004), which can be solved with standard convex optimisation solvers (e.g. SeDuMi, MOSEK, SDPT3). Specifically, we introduce a scalar slack variable a and solve

$$\begin{aligned} \min_{m\in\mathbb{R}^n, a\in\mathbb{R}} \quad & a \\ \text{s.t.} \quad & A_s(m) \preceq a I_n. \end{aligned} \tag{2.14}$$

If the optimal solution $a^* < 0$ then $A_s(m) \prec 0$ for some m, confirming that a trapping region exists. Conversely, if $a^* \geqslant 0$ then there are no coordinates under which the kinetic energy, $K_m(y) = (1/2)y^\top y$, verifies the existence of a trapping region for this system (Liao *et al.* 2025).

If a trapping region exists, we next refine the size of the trapping region. Instead of using the conservative estimate (2.13), we solve a quadratically constrained quadratic program (QCQP) to compute the exact minimal radius R_m^*:

$$(R_m^*)^2 = \max_y \; y^\top y \tag{2.15a}$$

$$\text{s. t. } d(m)^\top y + y^\top A_s(m)\, y \geqslant 0. \tag{2.15b}$$

Although this QCQP is non-convex, the dual function of its Lagrangian form admits a convex SDP dual problem (2.16) that can be solved to global optimality (Boyd & Vandenberghe 2004):

$$(R_m^*)^2 = \min_{\lambda\geqslant 0,\beta} \; \beta \tag{2.16a}$$

$$\text{s.t.} \begin{bmatrix} I_n + \lambda\, A_s(m) & \dfrac{\lambda}{2} d(m) \\ \dfrac{\lambda}{2} d(m)^\top & -\beta \end{bmatrix} \preceq 0. \tag{2.16b}$$

Here, λ is the primal-Lagrangian variable and β is a dual variable. Solving (2.16) yields R_m^*, the radius of the tightest certified Euclidean trapping ball $B(m, R_m^*)$ for a given centre m within the quadratic-energy certificate class. When solving (2.16), the shift m is treated as fixed (e.g. obtained from the feasibility problem (2.14) or from the stability-enforcement step discussed later in § 3 for a learned model). Jointly optimising over both the centre m and the trapping-ball radius would couple $A_S(m)$ and $d(m)$ with the SDP variables, resulting in a nonlinear (and generally non-convex) semidefinite program. We therefore do not attempt a joint optimisation in the present work, and instead adopt the sequential procedure: first certify feasibility by finding a shift m for which $A_S(m) \prec 0$ and then compute the globally optimal certified radius R_m^* for the given shift m. By leveraging this convex formulation, we avoid the conservatism of the original approach, enabling more precise and physically realistic bounds on the system's dynamics.

Using several numerical examples, we demonstrated in Liao *et al.* (2025) that both the trapping-region existence certificate and the optimal radius can be solved efficiently

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

with standard SDP solvers, yielding bounds often orders of magnitude tighter than prior analytical estimates. A discussion of computational complexity and scaling for the SDPs (2.14) and (2.16) can be found in Liao *et al.* (2025), where analysis of systems with state dimension $n \sim \mathcal{O}(100)$ were reported using a standard desktop computer. These convex formulations form the backbone of the present data-driven identification loop, ensuring that, for any choice of hyperparameters, we can both guarantee whether a lossless, globally bounded ROM exists and, if it does, compute its optimal trapping region in a single unified framework.

2.4. *Identification of sparse and lossless ROMs from data*

The ability to certifiably determine optimal trapping regions – as discussed in § 2.3 – can be leveraged within data-driven modelling frameworks, such as the operator inference framework (Peherstorfer & Willcox 2016; Goyal *et al.* 2024; Kramer *et al.* 2024) and the SINDy framework (Brunton *et al.* 2016; Loiseau & Brunton 2018). In this work, we demonstrate integration within the SINDy modelling framework, which is an optimisation-based regression algorithm commonly used for identifying sparse governing equations. The sparsity-promoting nature of SINDy enables the construction of more robust models despite errors in finite datasets, as promoting sparsity tends to prevent error propagation through coupled state variables.

In SINDy we first collect n_t 'snapshots' of the time evolution of the n-dimensional state x and its time derivative $\dot{x}$ into data matrices $X \in \mathbb{R}^{n_t \times n}$ and $\dot{X} \in \mathbb{R}^{n_t \times n}$, respectively. Each column of X corresponds to the time series of an individual state variable $x_i(t)$ sampled at discrete time points $t_1, t_2, \ldots, t_{n_t}$, and similarly for $\dot{X}$. In this work, we determine $\dot{X}$ from the data in X using an eighth-order central difference scheme. The data matrices X and $\dot{X}$ are then related by the dynamics, and SINDy provides a convenient way to express the right-hand side of (2.1) as the product of a candidate function library $\Theta(X)$ and a coefficient matrix Ξ:

$$\dot{X} = \Theta(X)\,\Xi, \qquad \Xi \in \mathbb{R}^{\left(1+n+\frac{n(n+1)}{2}\right)\times n}, \quad \Theta(X) \in \mathbb{R}^{n_t \times \left(1+n+\frac{n(n+1)}{2}\right)}. \tag{2.17}$$

The candidate function library $\Theta(X)$ is a data matrix consisting of constant (affine), linear and quadratic functions of the columns of X as needed to describe the flow. The candidate function library is constructed as

$$\Theta(X) = \begin{bmatrix} \vdots & \vdots & \vdots \\ 1 & X^{P_1} & X^{P_2} \\ \vdots & \vdots & \vdots \end{bmatrix}, \tag{2.18}$$

where the first column represents the affine terms, X^{P_1} contains linear functions of the data in X and X^{P_2} contains quadratic functions of the data in X. Note that we have chosen X^{P_2} to contain the unique quadratic monomials $x_j x_k$ with $j \leqslant k$ (a non-redundant representation). Equivalently, the columns of X^{P_2} are indexed by the upper-triangular index pairs (j, k) of a symmetric quadratic form. To correctly represent $x^\top Q^{(i)} x$ with symmetric $Q^{(i)}$, the off-diagonal monomials $x_j x_k$ for $j < k$ in X^{P_2} are scaled by a factor of 2, since each cross-term appears twice in the full expansion. The associated coefficient matrix Ξ is defined as

$$\Xi = \begin{pmatrix} c^\top \\ L \\ Q^\triangle \end{pmatrix}, \quad c \in \mathbb{R}^n, \qquad L \in \mathbb{R}^{n\times n}, \quad Q^\triangle \in \mathbb{R}^{\left(\frac{n(n+1)}{2}\right)\times n}, \tag{2.19}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

and includes the coefficients associated with the affine terms c, linear terms L and quadratic terms $Q^{\triangle}$.

The coefficient matrix Ξ in (2.17) is unknown, but can be determined via regression. The SINDy framework exploits the fact that only a few terms in $\Theta(X)$ will contribute to the dynamics of the flow, and so seeks to find a sparse coefficient matrix Ξ that best represents the dynamics. We note that even if the true underlying equations of motion are not sparse, promoting sparsity has the advantage of introducing model robustness to truncation errors and other factors. In addition to promoting sparsity in Ξ, we must ensure that the resulting model will possess quadratic terms that are lossless, which can be achieved by solving for a coefficient matrix Ξ that enforces $x^\top f(x) = 0$. To this end, we define a constrained least-squares regression problem with an L_1-regularisation (sparsity-promoting) term as (Loiseau & Brunton 2018)

$$\begin{aligned} \min_{\Xi} \quad & \left\| \Theta(X)\,\Xi - \dot{X} \right\|_2^2 + \delta \|\Xi\|_1 \\ \text{s.t.} \quad & C\,\Xi(:) = d. \end{aligned} \tag{2.20}$$

Here, the term $\delta\|\Xi\|_1$ promotes sparsity in Ξ, with the parameter δ controlling the degree to which sparsity is promoted. A larger δ promotes a sparser solution, while a smaller δ results in more active terms in Ξ; in general, the degree of sparsity with respect to δ is problem dependent and cannot be determined *a priori*. The linear constraint $C\,\Xi(:) = d$, where $\Xi(:) = \xi$ is the vectorised form of the matrix Ξ, ensures the identified system is lossless by encoding algebraic constraints on the quadratic entries of $\Xi(:)$ via the matrix C and vector d. For an explanation of how the constraint matrix C and the vector d are determined, refer to Appendix B.

By jointly enforcing sparsity and the lossless property, SINDy can be used to find a model that is physically consistent with governing incompressible Navier–Stokes equations. However, enforcing the nonlinear terms to be lossless is not sufficient to guarantee that resulting model trajectories will be long-time bounded. Thus, the resulting model may not be dynamically consistent with respect to the flow being modelled. For incompressible flows that admit a trapping region, matching the optimal trapping-region size provides a means of obtaining models that are (at least approximately) dynamically consistent. In the next section we describe how the convex optimisation problems associated with optimal trapping regions from § 2.3 can be integrated within the SINDy framework.

3. Optimally bounded nonlinear models from data

We now describe how to combine the convex optimisation methods described in 2.3 with the data-driven identification of a sparse, lossless and bounded quadratic ROM via SINDy, as described in § 2.4. To begin, we represent the system (2.1) in terms of a candidate function library and coefficient matrix, as in the SINDy formulation (2.17). We then modify the constrained least-squares problem (2.20) to simultaneously enforce boundedness, sparsity and losslessness in the identified dynamics.

Our objective is to find the coefficient matrix Ξ such that the resulting model is sparse, conserves energy via lossless quadratic terms and admits a provably optimal trapping region that guarantees long-time boundedness of model trajectories. In addition to solving for the coefficient matrix Ξ, we must also find the shift vector m. Recall from § 2.2 that by shifting the system to be centred at m such that $A_S(m)$ is negative definite, we ensure that the energy (2.9) decreases outside the trapping region. As discussed in § 2.3, we certify this condition in a single convex program by introducing a scalar slack variable

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

a to relax the strict negative-definiteness condition into a tractable inequality. Thus, for a given coordinate shift m, we seek the trapping region $B(m, R_m)$ centred about m that has the smallest radius. This is formulated as an optimisation problem that seeks to minimise the difference between the model prediction and observed data, while promoting a sparse solution for Ξ and enforcing system stability. The optimisation problem is defined as

$$\begin{aligned} \min_{\Xi,m,a} \quad & \|\Theta(X)\Xi - \dot{X}\|_2^2 + \delta\|\Xi\|_1 + a + \eta\,\|m\|_2 \\ \text{s.t.} \quad & C\Xi(:) = d \\ & A_S(m, \Xi) \preceq a I_n \\ & -\gamma \leqslant a. \end{aligned} \tag{3.1}$$

Just as in the SINDy equation (2.20), the first term of the objective function is the least-squares error between the predicted dynamics $\Theta(X)\Xi$ and the observed time derivatives $\dot{X}$, while the second term is an ℓ_1-regularisation term that promotes sparsity in Ξ. The third term, a, is the previously discussed slack variable for the trapping-region condition. The final term in the cost function $\eta\|m\|^2$ is added to regularise the optimisation to promote smaller shifts when multiple m satisfy the trapping condition, improving numerical conditioning. The constraint $C\Xi(:) = d$ ensures that the nonlinearity is lossless. The second constraint $A_S(m, \Xi) \preceq a I_n$ and the objective term that minimises a provide a boundedness certificate whenever the optimal slack satisfies $a^* < 0$; in this case, $A_S(m, \Xi) \prec 0$ and the system admits a globally monotonically attracting trapping ball. Note that the term $A_S(m, \Xi)$ has been modified from (2.8) to be a function of both m and Ξ. This notation is introduced because $A_S(m)$ is dependent on the linear term L and quadratic terms $Q^{(i)}$, which are now encoded in the decision variable Ξ. Finally, to prevent the objective from becoming unbounded from below, we impose the constraint $-\gamma \leqslant a$, which places a user-specified lower bound on the slack variable. Without this lower-bound constraint, we observed the optimisation to drive a to very negative values while selecting excessively large coordinate shifts m. These large shifts can render $A_S(m, \Xi)$ strongly negative definite, yet remain only weakly penalised through the regularisation term $\eta\|m\|_2$. Such solutions correspond to trapping certificates centred far outside the region explored by the training data and are not meaningful for the intended modelling task. The parameter γ therefore limits the enforced dissipation level and keeps the search in a physically relevant range.

This formulation enforces boundedness through a semidefinite constraint in (3.1), so any model retained after the hyperparameter sweep is accompanied by an explicit certificate (and models with $a^* \geqslant 0$ are discarded). In this sense, our approach differs from stability-promoting variants that rely on non-convex penalties or gradient-based stabilisation steps (Kaptanoglu *et al.* 2021).

3.1. *The trapping-SDP modelling framework*

The trapping-SDP procedure presented here solves for both the coefficient matrix Ξ and the shift vector m while promoting sparsity, boundedness and enforcing losslessness in the nonlinear term. To solve the problem using convex optimisation techniques, we consider two alternating steps: a regression step for the coefficient matrix Ξ and a stability-enforcement step for the shift vector m and a slack variable a. This is because the full optimisation problem (3.1) is not jointly convex in (Ξ, m, a): the semidefinite constraint depends on both the regression coefficients (through L and $Q^{(i)}$ encoded in Ξ) and the shift m. Solving (3.1) directly would therefore require a non-convex optimisation procedure. We instead adopt an alternating (block coordinate descent)

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

strategy that yields two convex subproblems: a regression step in Ξ with (m, a) fixed and a stability-enforcement step in (m, a) with Ξ fixed.

This alternating minimisation approach – sometimes referred to as block coordinate descent – is commonly used for problems where separate components (in this case, sparsity and stability) influence different parts of the optimisation problem. It simplifies each subproblem and accelerates convergence. The procedure proceeds in the following steps:

(i) variable initialisation,
(ii) regression,
(iii) stability enforcement,
(iv) iteration of steps (ii) and (iii),
(v) sparsity refinement.

This framework depends on two hyperparameters: the sparsity-promoting weight δ and the stability-enforcement term γ. We are interested in determining a model that best balances the desire for a small modelling error, a large degree of sparsity and a small trapping region. Therefore, we sweep over a range of values for δ and γ, repeating the modelling procedure described above for each parameter pair and then selecting the model that best balances the desired features.

For each candidate model resulting from the parameter sweep, we evaluate the performance using a validation-based procedure. Specifically, we first train the model on a finite trajectory from the training data. Then, we generate validation trajectories by time marching the identified model from a set of off-training initial conditions, ensuring the model is validated against data that was not seen during the training phase.

The model's accuracy is then assessed by computing the root-mean-square error (RMSE) between the predicted dynamics $x_i(t)$ and the validation data $x_{i,true}(t)$:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\big(x_i(t) - x_{i,true}(t)\big)^2}. \tag{3.2}$$

We compute the trapping region associated with each model and record its radius. As a measure of model sparsity, we use the sparsity ratio, i.e. the ratio of the number zero elements to the total number of elements in Ξ. The final model to be selected is the one that best balances a small prediction error, large sparsity ratio and a small trapping-region radius. In practice, we observe that models with the lowest prediction error often also possess the tightest or nearly tightest trapping regions. This multi-criterion model selection approach ensures that the final model is accurate, robust to numerical error and globally bounded.

We will now discuss each step of the procedure in greater detail.

3.1.1. *Initialisation*

It is necessary to first test whether a globally attracting trapping region can exist for the provided training data and hyperparameter γ. We therefore begin by computing an unconstrained least-squares estimate of the coefficient library $\Xi^{(0)}$, which provides an approximation of the dynamics in the chosen function library $\Theta(X)$:

$$\begin{aligned} \Xi^{(0)} \leftarrow \arg\min_{\Xi} \ & \frac{1}{2}\left\|\Theta(X)\,\Xi - \dot{X}\right\|_2^2, \\ \text{subject to} \quad & C\,\Xi(:) = d. \end{aligned} \tag{3.3}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

This initial fit ensures that the initial coefficient matrix $\Xi^{(0)}$ lies in a region of the parameter space that accurately reflects the underlying physics, rather than an arbitrary model.

Problem (3.3) is a convex equality-constrained least-squares program. In our implementation, the lossless constraints are homogeneous with $d = 0$ (Appendix B), so the feasible set $\{\Xi : C\,\Xi(:) = d\}$ is non-empty (i.e. $\Xi(:)$ is in the nullspace of C) and at least one global minimiser exists (Boyd & Vandenberghe 2004). The constraint matrix C is constructed to act only on the quadratic coefficient block of Ξ, leaving the affine and linear blocks unconstrained by $C\,\Xi(:) = d$. We therefore set $\delta = 0$ in the initialisation so that sparsity is not imposed before boundedness feasibility is assessed; sparsity promotion is introduced only in subsequent iterations via (3.6). Furthermore, solving a least-squares problem with the lossless constraints $C\,\Xi(:) = d$ preserves any known symmetry or conservation laws from the outset, preventing later iterations from chasing infeasible or physically inconsistent solutions. Finally, this initialisation provides a data-consistent starting point for the SDP step: since $\Xi^{(0)}$ already captures the principal dynamics, the subsequent searches for a trapping region via coordinate shift m and slack variable a are more likely to find a valid bounded solution, rather than failing due to an ill-posed initial guess. In practice, this regression-first strategy greatly accelerates convergence of the alternating optimisation and enhances the robustness of the final sparse, lossless and globally bounded ROM.

Once the coefficient matrix $\Xi^{(0)}$ is obtained, it is necessary to verify whether the current model admits a globally attracting trapping region and, if so, to compute an initial estimate of its centre. By the Schlegel–Noack trapping theorem, there exists a trapping region if and only if one can find a shift vector $m \in \mathbb{R}^n$ such that the shifted symmetric operator

$$A_S(m) \;=\; L_S\big(\Xi^{(0)}\big) \;-\; \sum_{i=1}^{n} m_i\, Q^{(i)}\big(\Xi^{(0)}\big) \tag{3.4}$$

is negative definite. As discussed in § 2.3, we certify this condition in a single convex program by using the slack variable a to relax the strict negative-definiteness condition into a tractable inequality:

$$\begin{aligned} \big(m^{(0)},\; a^{(0)}\big) \;\leftarrow\; \arg\min_{m,\,a} \big(a + \eta\,\|m\|_2\big), \\ \text{subject to} \quad A_S\big(m,\, \Xi^{(0)}\big) \preceq a\, I_n, \\ -\gamma \leqslant a. \end{aligned} \tag{3.5}$$

As in (3.1), the constraint $-\gamma \leqslant a$ prevents the SDP from being unbounded from below. Similarly, the regularisation term $\eta\|m\|^2$ biases the solution toward smaller shifts when multiple m satisfy the trapping condition. Upon convergence, a negative optimal value $a^{(0)} < 0$ provides a rigorous certificate that a globally attracting trapping region exists for the initial ROM, allowing the initial estimates $(m^{(0)}, a^{(0)})$ to then seed the alternating regression and stability-enforcement iterations that jointly refine both Ξ and the trapping-region parameters. Conversely, in the case that $a^{(0)} \geqslant 0$, we can definitively conclude that a globally attracting trapping region does not exist for the current model coefficients, and so can discard the model. We emphasise that $a^{(0)} \geqslant 0$ does not rule out the existence of bounded dynamics in general; rather, it indicates that the initial (lossless) least-squares fit cannot be certified as globally bounded within the adopted quadratic-energy certificate class. Because the subsequent regression step (3.6) enforces boundedness by requiring feasibility with respect to the previous iterate's certificate, we require an initial feasible certificate (i.e. $a^{(0)} < 0$) to start the alternating loop. If

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

no hyperparameter choice yields $a^{(0)} < 0$ then an alternative library, model order or boundedness certificate class is required. This two-step initialisation is both theoretically grounded and practically effective: the least-squares fit captures the dominant dynamics from data and the SDP then determines whether global boundedness is attainable. Together, they provide a starting point for seeking a model with an optimal trapping region.

3.1.2. *Iterative regression and stability enforcement*

We next proceed with an alternating block coordinate descent loop, updating the coefficients and the stability certificate in turn until both have converged. For the regression step, we solve the following convex optimisation problem to update Ξ at each iteration k, with $m^{(k-1)}$ and $a^{(k-1)}$ fixed:

$$\begin{aligned} \Xi^{(k)} \leftarrow \arg\min_{\Xi} \quad & \left\| \Theta(X)\Xi - \dot{X} \right\|_2^2 + \delta \|\Xi\|_1 \\ \text{s.t.} \quad & C\Xi(:) = d \\ & A_S\left(m^{(k-1)}, \Xi\right) \preceq a^{(k-1)} I_n. \end{aligned} \tag{3.6}$$

Here, the ℓ_1 penalty $\delta\|\Xi\|_1$ promotes sparsity in the coefficient matrix Ξ, while the semidefinite constraint enforces that the current linear and quadratic operators admit a trapping region when shifted by $m^{(k-1)}$. Because this is a convex program, we can obtain the updated estimate for $\Xi^{(k)}$ efficiently. Next, holding $\Xi^{(k)}$ fixed, we refine the shift vector m and slack variable a by solving

$$\begin{aligned} m^{(k)}, a^{(k)} \leftarrow \arg\min_{m,a} \quad & a + \eta \|m\|_2 \\ \text{s.t.} \quad & A_S\left(m, \Xi^{(k)}\right) \preceq a I_n \\ & -\gamma \leqslant a. \end{aligned} \tag{3.7}$$

This SDP updates the trapping-region centre m and slack variable a so that the shifted symmetric operator $A_S(m, \Xi^{(k)})$ remains negative definite, thereby enforcing global boundedness for the current model. Again, we use a small η so that the $\eta\|m\|_2$ term regularises the shift so as to avoid ill-conditioned solutions. We iterate between the regression and stability-enforcement steps until convergence is achieved, based on the following criteria:

$$\left|\Xi^{(k)} - \Xi^{(k-1)}\right| \leqslant \epsilon_\Xi \quad \text{and} \quad \left|m^{(k)} - m^{(k-1)}\right| \leqslant \epsilon_m. \tag{3.8}$$

At convergence, the trapping-region condition is checked. If $a^{(k)} \geqslant 0$, the identified model is not guaranteed to be bounded and, thus, removed as a candidate. Just as in the initialisation, infeasibility at this step indicates that boundedness cannot be certified within the chosen model structure and certificate class, not that the underlying system is necessarily unbounded. If $a^{(k)} < 0$, we obtain a model that simultaneously fits the data, is sparse as prescribed by the hyperparameter δ and provably admits an optimal trapping region. Since boundedness certification is complementary to predictive accuracy, our model selection balances the certified bound size with validation error and sparsity. This alternating scheme leverages the convexity of each subproblem to ensure efficient and reliable computation of the final ROM.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

3.1.3. *Convergence and refinement*

We now briefly discuss the convergence properties of the algorithm and the final refinement step. Let

$$J\left(\Xi^{(k)}, m^{(k)}, a^{(k)}\right) = \frac{1}{2}\left\|\Theta(X)\,\Xi^{(k)} - \dot{X}\right\|_2^2 + \delta\left\|\Xi^{(k)}\right\|_1 + a^{(k)} + \eta\left\|m^{(k)}\right\|_2 \tag{3.9}$$

denote the objective function at iteration k. We note that the objective is bounded from below by $-\gamma$ due to the constraint $a \geqslant -\gamma$. At the initial iterate ($k = 0$), the objective function is $J(\Xi^{(0)}, m^{(0)}, a^{(0)})$ by virtue of initialisation. At each subsequent iterate k,

$$J\left(\Xi^{(k-1)}, m^{(k-1)}, a^{(k-1)}\right) \geqslant J\left(\Xi^{(k)}, m^{(k-1)}, a^{(k-1)}\right) \geqslant J\left(\Xi^{(k)}, m^{(k)}, a^{(k)}\right), \tag{3.10}$$

where the first inequality holds because $\Xi^{(k)}$ minimises $J(\Xi, m^{(k-1)}, a^{(k-1)})$ over a feasible set that contains $\Xi^{(k-1)}$, and the second inequality holds because $(m^{(k)}, a^{(k)})$ minimises $J(\Xi^{(k)}, m, a)$ over a feasible set that contains $(m^{(k-1)}, a^{(k-1)})$.

After iterating through the regression and stability steps, we check for convergence based on the changes in Ξ and m. The algorithm terminates when both updates fall below the prescribed tolerances ϵ_Ξ and ϵ_m. From an algorithmic standpoint, while each subproblem in the alternating procedure is convex, the combined problem (3.1) is not jointly convex in (Ξ, m, a), so convergence is to a stationary point that can depend on initialisation and hyperparameters; this behaviour is typical of block coordinate descent methods for non-convex objectives (Tseng 2001). If convergence is achieved, we refine the solution by fixing the sparsity pattern and forcing small terms in Ξ to zero. Let Ξ^K be the converged coefficient vector and define a tolerance such that any $|\Xi_i^K| < \epsilon_x$ is considered small. We construct the matrices C_{ref} and d_{ref} as

$$C_{ref}(i,\,j) = \begin{cases} 1 & \text{if } \left|\Xi_j^K\right| < \epsilon_x, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad d_{ref} = \mathbf{0}. \tag{3.11}$$

The augmented constraint matrices are

$$\tilde{C} = \begin{pmatrix} C \\ C_{ref} \end{pmatrix}, \quad \tilde{d} = \begin{pmatrix} d \\ d_{ref} \end{pmatrix}. \tag{3.12}$$

Using the optimal solutions for the shift vector m^* and the slack variable a^*, obtained in the last iteration of (3.7), we then solve the optimisation problem one final time with these updated constraints to ensure the small terms are forced to zero:

$$\begin{aligned} \Xi^* \leftarrow \arg\min_{\Xi} \quad & \left\|\Theta(X)\Xi - \dot{X}\right\|_2^2 + \delta\|\Xi\|_1 \\ \text{s.t.} \quad & \tilde{C}\Xi = \tilde{d} \\ & A_S(m^*,\,\Xi) \preceq a^* I_n. \end{aligned} \tag{3.13}$$

This step finalises the sparse and bounded model by eliminating negligible coefficients. We note that the refinement step could increase the cost

$$J\left(\Xi^{(k-1)}, m^{(k-1)}, a^{(k-1)}\right) \leqslant J^*(\Xi^*, m^*, a^*) \tag{3.14}$$

as more constraints are imposed during this step. However, since we terminate when both $\|\Xi^{(k)} - \Xi^{(k-1)}\|$ and $\|(m^{(k)}, a^{(k)}) - (m^{(k-1)}, a^{(k-1)})\|$ fall below user-specified tolerances, we ensure a self-consistent, locally optimal solution.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

4. Numerical benchmark examples

In this section we provide two numerical benchmark examples to illustrate the proposed method on systems with known dynamics: (i) a Lorenz dynamical system (Lorenz 1963), and (ii) a ROM of a sinusoidally forced shear flow (Moehlis *et al.* 2004). Application to an unsteady nonlinear aerodynamics problem follows in § 5.

4.1. *Lorenz system*

To illustrate our approach, we first examine the Lorenz dynamical system (Lorenz 1963)

$$\frac{dx_1}{dt} = -\sigma x_1 + \sigma x_2, \tag{4.1a}$$

$$\frac{dx_2}{dt} = \rho x_1 - x_2 - x_1 x_3, \tag{4.1b}$$

$$\frac{dx_3}{dt} = -\alpha x_3 + x_1 x_2, \tag{4.1c}$$

with parameters $\sigma = 10$, $\rho = 28$ and $\alpha = 8/3 \approx 2.667$. Although the Lorenz system exhibits chaotic behaviour and lacks stable equilibrium points, it has been shown to be bounded (Schlegel & Noack 2015). In our analysis, we verify that by applying a translation m, the symmetric part of the linear operator, $A_s(m)$, becomes negative definite. According to the theorem discussed in § 2.2, this implies the existence of a trapping region $B(m, R_m)$ with an estimated radius $R_m = 101.33$. We note that this radius slightly differs from that reported in Schlegel & Noack (2015), due to a correction noted by Kaptanoglu *et al.* (2021).

In our previous work (Liao *et al.* 2025), we used our trapping-SDP method to determine a trapping region for the Lorenz system. There, we obtained a translation $m = [0,\ 0,\ 38]^\top$ and an optimal trapping-region radius of $R_m^* = 39.25$, which is significantly smaller than the conservative bound of $R_m = 101.33$. It is important to note that these results were achieved with full knowledge of the system dynamics. In contrast, our current objective is to employ our data-driven algorithm to recover the dynamic equations from a sample trajectory of the Lorenz system without prior knowledge of the governing equations, and to obtain a trapping-region size close to the optimal value determined previously.

To this end, we integrated the Lorenz equations (4.1) forward in time over 5000 time steps from a random initial condition. We then applied the algorithm described in § 3, sweeping over a range of sparsity (δ) and stability (γ) parameters, ultimately choosing a model that balanced accuracy and trapping-region size (as shown in figure 1). Our method successfully recovered the sparsity pattern of the original Lorenz system with parameters $\sigma = 9.9903$, $\rho = 27.9684$ and $\alpha = 2.6661$. The trapping region for the identified model was $R_m^* = 39.3144$, closely matching the optimal value from our earlier work (Liao *et al.* 2025). These results are summarised in table 1.

To validate the results, we integrated the discovered model from a set of initial conditions not used for training and compared them to the true dynamics. These results are summarised in figure 2, where we also compare the trapping region identified using our method with that of Schlegel and Noack. Note that the energy of the shifted system $K_m(x(t))$ decays monotonically towards the trapping region.

4.2. *Sinusoidally forced shear flow*

Next, we consider the nine-state Galerkin model of a sinusoidally forced shear flow introduced by Moehlis *et al.* (2004), which has been used in a number of previous works

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

Parameter	σ	ρ	α	Trapping-region size R_m^*
Lorenz	10	28	$8/3 \approx 2.667$	39.25
Model result	9.9903	27.9684	2.6661	39.3144

Table 1. Tabulated comparison of the true Lorenz parameters used to generate the training trajectory with the parameters obtained using our method (labelled as 'model result'). The table also compares the optimal trapping region of the true Lorenz system with the trapping region of the model.

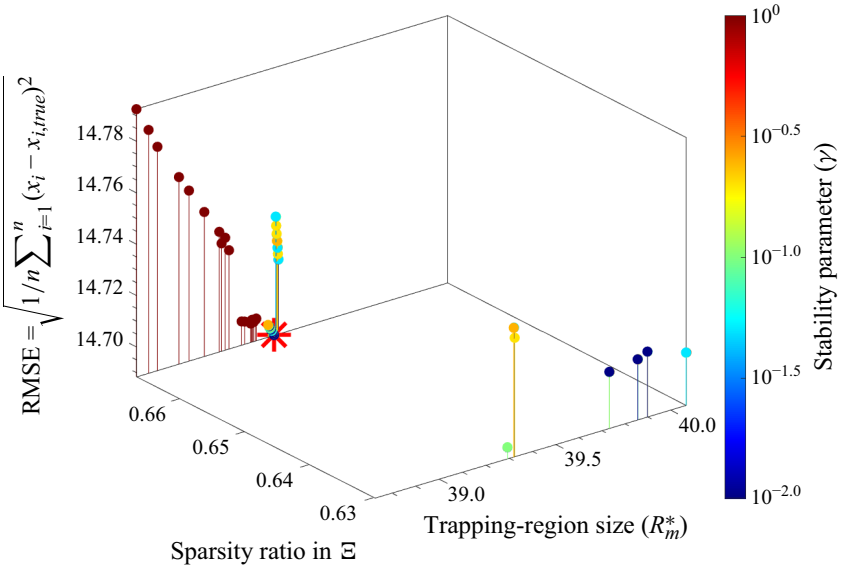


Figure 1. Model selection for the three-state Lorenz ROM. Each marker denotes a candidate model obtained by sweeping the sparsity parameter δ (horizontal axis: sparsity ratio in Ξ) across two levels and varying the stability weight γ (colour scale). The trapping-region size R_m^* (depth axis) and the RMSE of the state predictions (vertical axis) are shown for each model. Because the Lorenz system has only three states, only two sparsity patterns arise, producing two vertical 'stems' of points. The denser stem (lower sparsity) yields a larger RMSE and looser bounds, while the sparser stem (higher sparsity) achieves both lower error and a tighter trapping region. The optimal model, indicated by the red star, corresponds to the higher-sparsity solution with the smallest RMSE and trapping region.

to demonstrate nonlinear stability analysis methods (see, e.g. Goulart & Chernyshenko 2012; Liu & Gayme 2020; Kalur *et al.* 2022; Buzhardt & Graham 2025; Enayati Kafshgarkolaei & Hemati 2025; Heide & Hemati 2025). The flow consists of a sinusoidal body force applied between two free-slip, counter-moving walls. The reduced-order Galerkin model is based on a projection of the dynamics onto nine Fourier modes, each capturing a physically important feature of the flow dynamics. For more details on the model and its physical significance, the reader is referred to the original source (Moehlis *et al.* 2004).

The nine-state model has the form

$$\dot{\tilde{x}} = f(\tilde{x}; Re) = \tilde{L}(Re)\tilde{x} + Q(\tilde{x})\tilde{x}, \tag{4.2}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

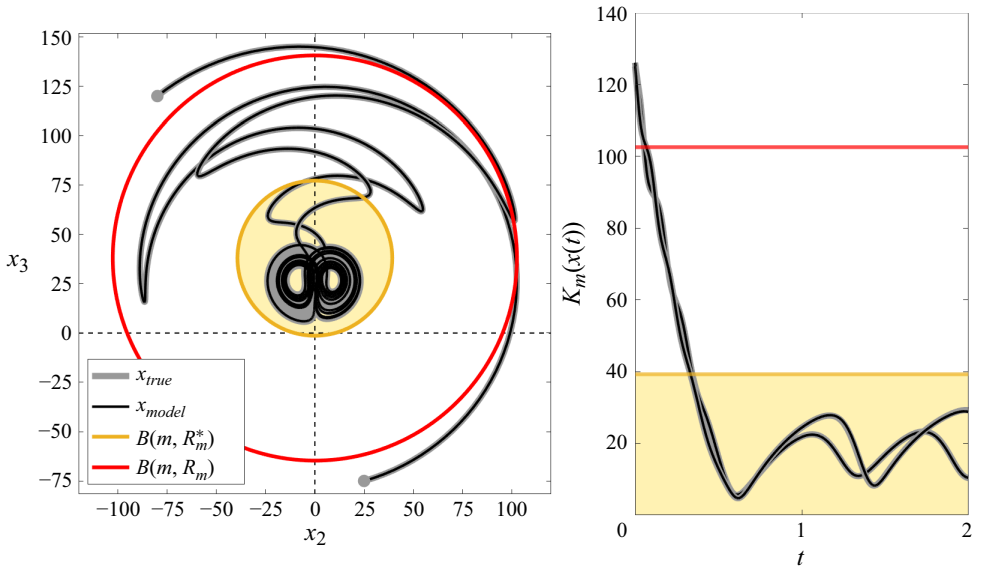


Figure 2. Trapping-region comparison. Here $B(m, R_m^*)$ shows the trapping region computed via our method and $B(m, R_m)$ shows the trapping region computed via the Schlegel and Noack method.

where the state vector $\tilde{x} = \tilde{x}(t) \in \mathbb{R}^9$ consists of nine modal coefficients that govern the temporal evolution of the flow. The quadratic term is defined as

$$Q(\tilde{x})\,\tilde{x} \;=\; \begin{bmatrix} \tilde{x}^\top Q^{(1)} \tilde{x} \\ \vdots \\ \tilde{x}^\top Q^{(9)} \tilde{x} \end{bmatrix}, \tag{4.3}$$

where each $Q^{(i)} \in \mathbb{R}^{9\times 9}$ is symmetric and $\tilde{x}^\top Q(\tilde{x})\,\tilde{x} \equiv 0$, i.e. the quadratic nonlinearity is lossless. The linear operator $\tilde{L}(Re)$ is Hurwitz and is parameterised by the Reynolds number $Re > 0$. In the original formulation the non-trivial equilibrium sits at

$$\tilde{x} = [\,1,\ 0, \ldots, 0\,]^\top, \quad \dot{\tilde{x}} = 0. \tag{4.4}$$

To recentre this fixed point at the origin, we introduce the shift

$$x = \tilde{x} - c, \qquad c = [\,1,\ 0, \ldots, 0\,]^\top, \tag{4.5}$$

so that $\dot{x} = 0$ at $x = 0$. This translation modifies the linear operator: if $\widetilde{L}(Re)$ is the original linear term then

$$L(Re) \;=\; \widetilde{L}(Re) \;+\; W, \tag{4.6}$$

where $W \in \mathbb{R}^{9\times 9}$ is chosen to absorb the cross-terms introduced by the shift, namely

$$W\,\tilde{x} \;=\; Q(\tilde{x})\,c \;+\; Q(c)\,\tilde{x}, \tag{4.7}$$

with $Q(\cdot)$ being the collection of symmetric quadratic operators.

We apply our modelling framework for a Reynolds number of $Re = 200$ to provide a stringent test of our method's ability to recover both the correct governing coefficients and a certifiably tight trapping region. As noted by Moehlis *et al.* (2004), the case of $Re = 200$ results in strongly nonlinear interactions with intermittent bursts. The strong energy

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

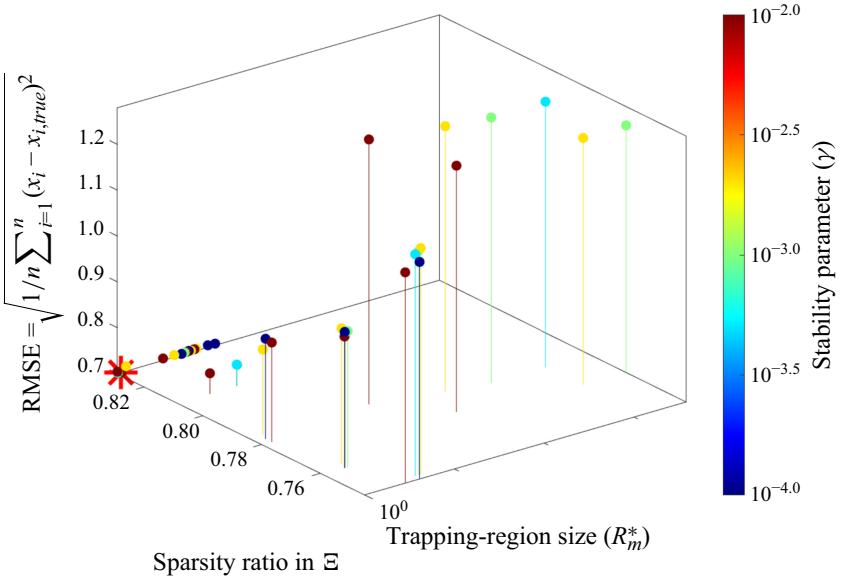


Figure 3. Model selection for the nine-state sinusoidal shear-flow ROM. Each marker represents a candidate model obtained by sweeping the sparsity parameter δ (horizontal axis: sparsity ratio in Ξ) and the stability weight γ (colour scale) through our alternating regression–SDP procedure. The trapping-region size R_m^* (log scale, depth axis) and the RMSE of the POD coefficient predictions (vertical axis) are plotted for each model. For each γ (colour), varying δ yields a vertical 'stem' of points: sparser models (higher δ) lie to the left, while denser models lie to the right. Note that two distinct clusters emerge, one with small trapping regions but higher RMSE (lower γ) and another with larger regions but lower RMSE (higher γ). The optimal model, marked by the red star, strikes the best compromise, achieving both low error and a tight trapping region.

exchanges and modal interactions in this flow make obtaining accurate and bounded ROMs from data challenging.

The model was initially trained on three concatenated simulations, each generated from random initial conditions in $[-3, 3]$ and evolved for 1000 time steps over $t \in [0, 1000]$; this range was chosen to ensure return to baseline within the simulation window. The model was then validated on a separate set of initial conditions to assess its ability to generalise beyond the training regime. A sweep over γ and δ was performed to generate a set of candidate models, and the model that best balanced a low error while maintaining a tight trapping region was selected for each case (see figure 3).

To determine the accuracy of the selected model, we first compared the errors between the modelled and true coefficients. The errors were computed via the Frobenius norm of the difference between the true and modelled coefficient matrices, normalised by the Frobenius norm of the true matrices. This resulted in a linear coefficient error of $\|L^* - L\|_F / \|L\|_F = 0.0165\,\%$ and a quadratic coefficient error of $\|Q^* - Q\|_F / \|Q\|_F = 0.027\,\%$. The errors are all of the order of a hundredth of a percent, thus indicating that both the linear and quadratic model coefficients are accurately recovered. Next, we compared the optimal trapping-region radius for our model (R_m^*) with that of the true system ($R_{m,true}^*$), yielding an error of $\|R_m^* - R_{m,true}^*\|_F / \|R_{m,true}^*\|_F = 0.001\,\%$. Furthermore, the shift vector $m^* = [-0.5, 0, 0, 0, 0, 0, 0, 0, -0.5]^\top$ of the recovered model was identical to that of the true system.

The trapping-SDP modelling approach correctly recovered the sparsity pattern of the true coefficient matrices. Figure 4 shows a comparison of the modelled coefficient matrices Ξ^* and the true coefficient matrices Ξ, where colour maps indicate active

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

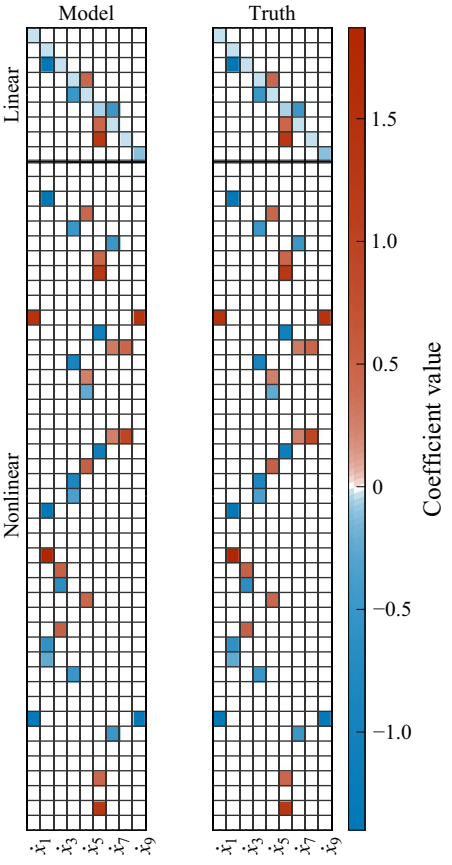


Figure 4. Coefficient matrices Ξ, for the $Re = 200$ sinusoidal shear flow, with modelled and true coefficients on the left and right, respectively. The sparsity pattern of the modelled and true coefficients is identical. The off-diagonal linear terms are a result of the shift W to ensure that the system is centred at the origin.

terms and inactive terms are rendered in white. The sparsity pattern of the modelled coefficients exactly matches that of the true coefficients, and the magnitudes match exactly. Additionally, figure 5 shows the time evolution of the system's energy along with the identified trapping region. The trajectories initiated from conditions outside the trapping region exhibit a monotonic decay until they enter the region, after which the energy remains bounded. The energy curves from the true model and the reconstructed model match closely.

5. Low-order model of unsteady separation over an aerofoil

Here, we apply our framework to obtain low-order models of unsteady separation over a NACA 65(1)-412 aerofoil at $Re = 20\,000$ (with respect to the chord) and an angle of attack of $\alpha = 4^\circ$ (see figure 6). Accurately modelling an aerofoil's baseline (uncontrolled) flow is critical for developing effective flow control strategies. Reduced-order models that capture the essential flow physics of baseline flows while simplifying model complexity are especially valuable for controller design (Brunton & Noack 2015; Rowley & Dawson 2017). Without a faithful baseline flow model, control laws may be ineffective or even destabilising.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

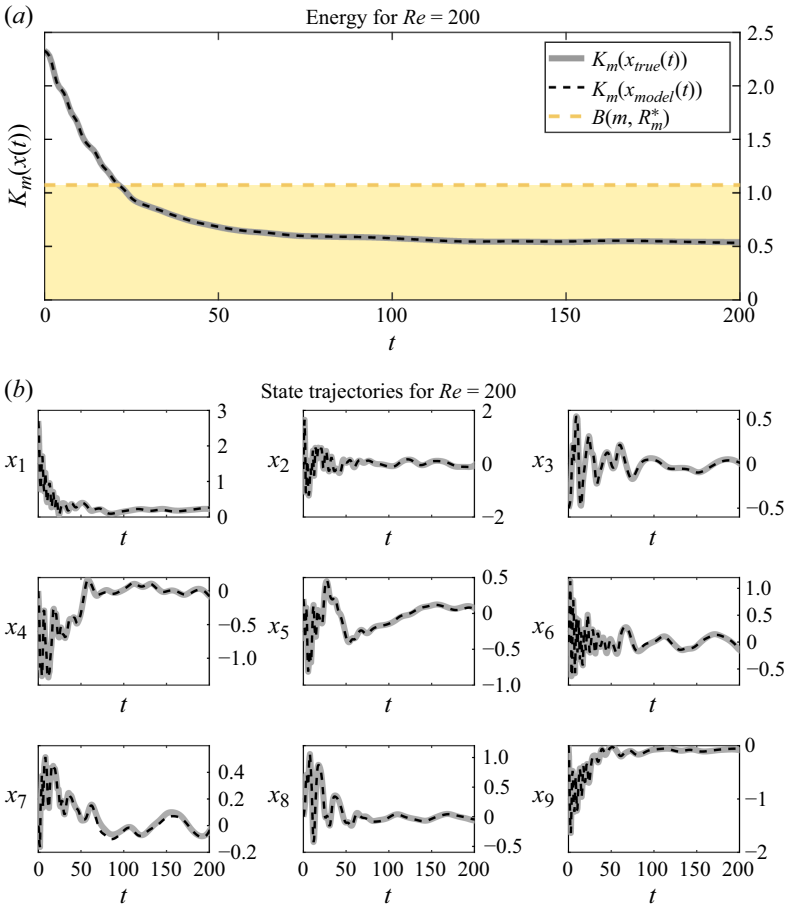


Figure 5. The energy and state validation trajectories for $Re = 200$. The grey indicates the true trajectory and the black dashed line indicates the trajectory obtained from integrating the modelled Ξ from off-training initial conditions. Note that the energy decays monotonically towards the trapping region. (*a*) Energy for $Re = 200$, (*b*) state trajectories for $Re = 200$.

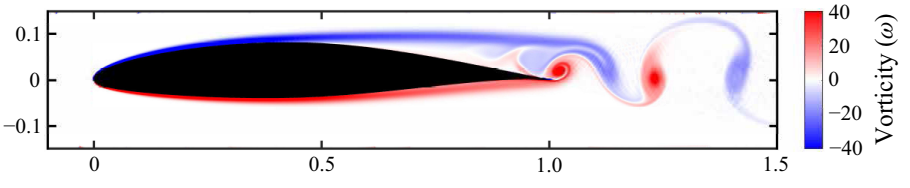


Figure 6. Baseline flow over a NACA 65(1)-412 aerofoil at $Re = 20\,000$ and $\alpha = 4°$.

5.1. *Direct numerical simulations and data curation*

For model training, we conducted DNS of the baseline flow. The flow considered is nominally incompressible; yet, the moderate Reynolds number and the strong separation bubble make it a challenging test for reduced-order modelling: the flow exhibits a broad spectrum of energetic scales, an unsteady separation region and a self-sustained oscillation (lift 'buzz') whose phase must be captured accurately (Tank *et al.* 2021). We also conducted DNS of the flow response to localised pulse forcing in order to provide a rich

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

dataset for off-training model validation. Localised forcing was applied at six chordwise locations on the suction side of the aerofoil,

$$z_x/c \in \{0.1,\ 0.2,\ 0.3,\ 0.4,\ 0.45,\ 0.5\}, \tag{5.1}$$

and at six discrete times corresponding to fractions of the natural lift-oscillation period,

$$T_{pulse} \in \left\{0,\ T_f/6,\ 2T_f/6,\ 3T_f/6,\ 4T_f/6,\ 5T_f/6\right\}, \tag{5.2}$$

yielding a total of 36 validation cases. Each pulse is an instantaneous, zero-mass, Gaussian-shaped forcing whose principal axis is oriented normal to the surface. The validation cases are initialised after the pulse-induced transient has decayed for model identification and testing of trapping-region performance. Specific details on the DNS are reported in Appendix C.

Prior to applying our modelling framework, we curate the DNS data by restricting the spatial domain to focus on the separation bubble and near-wake region (see figure 6). This restriction concentrates the modelling effort on the suction-side separation dynamics above the aerofoil and the immediate near wake, which are the primary regions of interest for reduced-order modelling and control, while reducing the number of spatial degrees of freedom. The full computational domain and numerical details are reported in C. The resulting spatial domain consists of $r = 48\,549$ spatial grid points and 550 temporal snapshots.

Next, we decompose the velocity field $\boldsymbol{u}(z, t)$ into a temporal-mean component $\bar{\boldsymbol{u}}(z)$ and a fluctuating component $\boldsymbol{u}'(z)$. Thus,

$$\boldsymbol{u}(z, t) = \bar{\boldsymbol{u}}(z) + \sum_{i=1}^{n} \boldsymbol{\phi}_i(z) x_i(t), \tag{5.3}$$

where we have expanded the fluctuations in a basis of POD modes $\boldsymbol{\phi}_i(z) \in \mathbb{R}^{2r}$ determined from snapshots of $\boldsymbol{u}'(z, t)$. In this work, we retain the leading $n = 6$ POD modes (see figure 7), which capture 80 % of the fluctuating kinetic energy in the baseline flow. Thus, the state vector to be used within the modelling framework will correspond to the six POD modal coefficients x_i. Note that the leading six POD modes are dominated by wake structures and that the first four modes in particular retain much of the larger-scale dynamics within the separation region above the aerofoil. The higher-order POD modes that have been excluded from the modelling exhibit finer-scale spatial structures; as such, we expect that fine-scale features will not be fully captured in the flow reconstruction based on (5.3).

Finally, we perform a time-interpolation of the POD coefficients to refine the time resolution ($\Delta t = 0.002$) and obtain POD coefficient time-derivative information from an eighth-order central differencing scheme.

This post-processing is conducted over a shorter time horizon than the original DNS, amounting to 14.5 % of the DNS simulation horizon. The total number of snapshots on this finer grid available for model training is then $n_t = 400$, from which snapshot data matrices X and $\dot{X}$ in (2.17) are constructed.

5.2. *Aerofoil model results*

We apply our modelling framework to the snapshot data matrices of POD coefficients, constructed as described above. In doing so, we swept over the sparsity-promoting parameter δ and the stability regularisation weight γ to generate a family of candidate models, each constrained to admit a monotonically attracting trapping region. From this

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

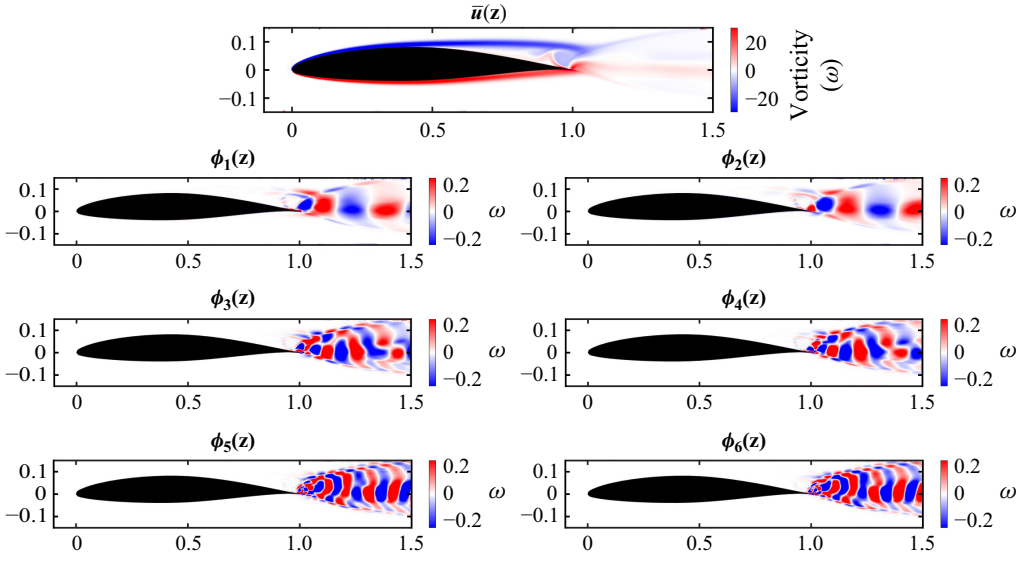


Figure 7. Mean flow $\bar{u}(z)$ and leading six POD modes $\phi_i(z)$ used in modelling flow over a NACA 65(1)-412 at $Re = 20\,000$ and $\alpha = 4°$, visualised via the associated vorticity field.

candidate pool, we selected the model that best balanced the RMSE of the modelled coefficients with respect to validation data and the size of the trapping region $B(m, R_m^*)$.

Figure 8 visualises this model selection process. We plot the RMS prediction error versus sparsity and trapping-region size. Models with larger trapping regions exhibit a higher prediction error, as do overly sparse models. These trends suggest that an optimal trade-off exists: although sparsity is desirable, some small but important coupling terms must be retained to accurately capture the dynamics. The best-performing model was 22 % sparse (i.e. 22 % of the model coefficients were set to zero) and yielded a minimal trapping-region radius of $R_m^* = 1\,452$. Note that application of the Schlegel and Noack method to the same model yields a trapping-region radius estimate of $R_m = 1\,930$, which is roughly 32 % more conservative than the optimal estimate. We report coefficients for the selected model in Appendix D.

In figure 9 we present the learned coefficient structure of the identified six-state model. Several features stand out. First, unlike the Lorenz and nine-state models, which omitted affine terms, the six-state ROM of the aerofoil benefits significantly from including affine terms in the candidate library. In practice, allowing each equation to include a constant bias term improved fit and boundedness: the learned affine shifts in $\dot{x}_1$, $\dot{x}_5$ and $\dot{x}_6$ captures a persistent mean-flow deviation or bias that cannot be represented by purely linear or quadratic terms. This highlights the importance of incorporating affine modes when modelling flows with non-zero equilibrium offsets, ensuring the ROM can adjust its baseline state as needed for accurate long-term predictions. Second, strong linear couplings are evident in the equations for $\dot{x}_1$, $\dot{x}_2$, $\dot{x}_3$ and $\dot{x}_5$. In contrast, the $\dot{x}_4$ equation contains no affine components and only a weak linear coupling with POD mode 5; its evolution is governed mainly by quadratic terms, underscoring the nonlinear nature of that mode. Finally, mode 5 not only carries its affine shift and a strong linear coupling with POD mode 6, but also features significant quadratic coefficients across multiple modal pairs. This indicates that, beyond its linear and constant contributions, mode 5 is subject to rich nonlinear modal interactions. While the model is sparse, the retained coefficients

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

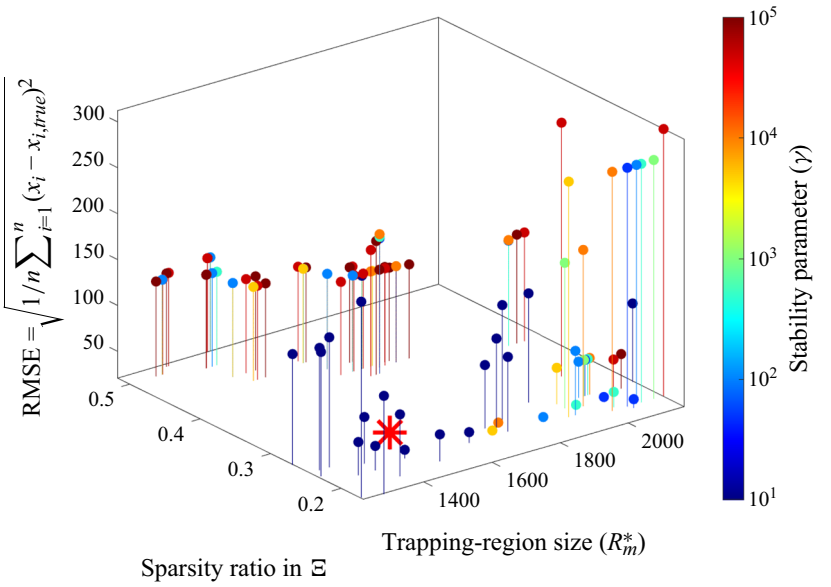


Figure 8. Model selection for the six-state aerofoil ROM. Each marker represents a candidate model obtained by sweeping the sparsity parameter δ (horizontal axis: sparsity ratio in Ξ) and the stability weight γ (colour scale from dark blue, $\gamma = 10^1$, to dark red, $\gamma = 10^5$). The trapping-region size R_m^* (depth axis) and the RMSE of the POD coefficient predictions (vertical axis) are shown for each model. Compared with the three- and nine-state cases, the aerofoil data yields a broader spread in both RMSE and trapping-region size, reflecting the increased system complexity. Although many models achieve similarly low RMSE (clustered near the front of the plot), they differ substantially in their certified trapping region. We selected the red-star model by first identifying the five models with the smallest RMSE and then choosing the one among them with the tightest bounding region. Importantly, all low-error candidates behave well under time integration, demonstrating robustness across hyperparameter choices.

(shown in grey) suggest meaningful nonlinear interactions, despite the fact that the model was trained exclusively on the unactuated baseline flow.

To assess generalisation within the baseline limit cycle, we validated the ROM on three off-training initial conditions drawn from phase-shifted points along the same unforced trajectory. This tests whether the ROM correctly reproduces the oscillatory dynamics over long time horizons. Results for coefficients x_1, x_3 and x_5 are shown in figure 10. The ROM (black dashed line) closely tracks the DNS (grey solid line) across all cases. A small phase drift is observed as time progresses, likely due to the exclusion of higher-order POD modes that may encode weak phase-coupling dynamics lost during truncation. Importantly, the model preserves the amplitude and qualitative features of the oscillation, suggesting robust internal limit cycle dynamics.

Although the model was trained only on baseline flow data, we further validated it on perturbed (pulsed) flow simulations. In these cases, the ROM was initialised with states taken after the actuation-induced transient had largely decayed, ensuring the validation remains within the ROM's attractor manifold. This isolates model fidelity to intrinsic dynamics rather than evaluating its response to exogenous forcing (which it was not trained to reproduce).

Figure 11 shows the selected validation windows for three actuation times at z_x/c, with the initial state (post-transient) marked by a red star. Although some residual unsteadiness remains, this set-up allows testing whether the ROM remains bounded and predictive from non-training states.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

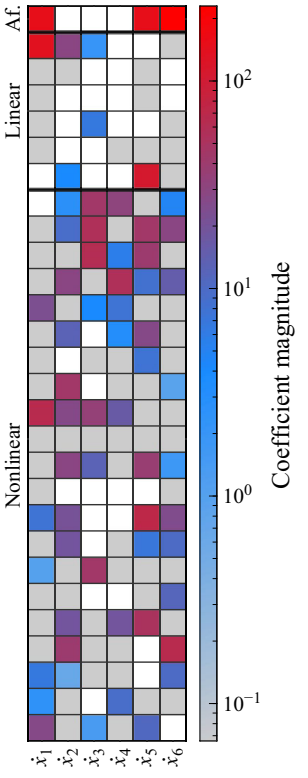


Figure 9. Coefficient magnitudes of the six-state ROM for the aerofoil flow. Linear, affine and quadratic terms are shown in separate blocks; colours indicate log-scaled coefficient magnitudes and zero (sparse) entries are rendered in white.

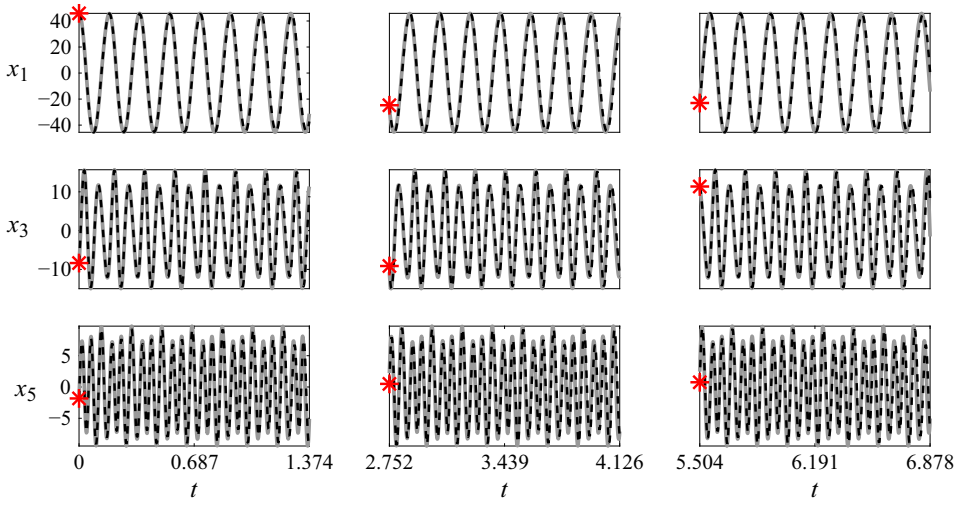


Figure 10. Validation of the six-state ROM on off-training initial conditions along the baseline limit cycle. Solid grey curves show the true POD coefficients, while black dashed curves show the ROM predictions. Each column corresponds to a different phase-shifted initialisation (marked by red stars). Displayed are modes x_1, x_3 and x_5; the remaining modes (x_2, x_4, x_6) exhibit similar agreement and are omitted for clarity.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

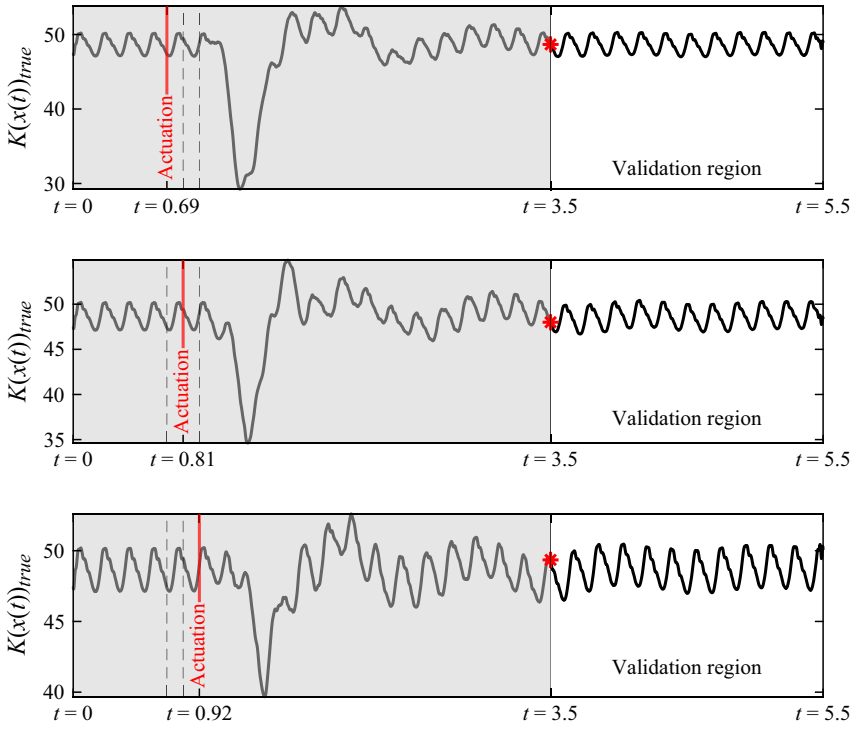


Figure 11. Time series of the true kinetic energy $K(x(t))_{true}$ for pulsed flow at $x/c = 0.1$, showing three sample actuation times: $t = 0\,T_f/6$, $t = 2\,T_f/6$ and $t = 4\,T_f/6$. In each subplot, the grey-shaded interval denotes the transient phase, the solid red line marks the actuation instant and the vertical dashed lines show the three other pulse timings for comparison. The black curve to the right of the shaded region shows the data used for validation, initialised at $t = 3.5$, indicated by the red star.

Figure 12 displays the energy evolution and trapping regions for the selected initialisations. In all cases, the ROM remains safely within its trapping region for the entire simulation horizon. However, the trapping-region bounds are much larger than in the previous examples. Several factors likely contribute to this comparatively large trapping region in the six-state aerofoil model. First, truncating the POD expansion to just six modes necessarily omits higher-order interactions that would otherwise strengthen the negative-definite component of the dynamics; without these modes, the identified quadratic operators must compensate, broadening the worst-case energy estimate. Second, the finite POD basis itself may not perfectly orthogonalise all nonlinear couplings, introducing spurious cross-terms that further inflate the radius required to cover every trajectory. Finally, the inclusion of a non-zero affine shift, most prominently in the $\dot{x}_5$ equation, displaces the centre m of the trapping ball away from the origin, which increases the minimal radius even though the actual energy peaks remain small. Nonetheless, this bound still certifies global boundedness: throughout every test case, the six-state ROM remains inside its invariant ball and faithfully reproduces the baseline flow amplitudes. We also note that similar to the unforced case, phase drift increases with time, especially evident in the second and third scenarios. For actuation cases at $t = 0T_f/6$ and $t = 4T_f/6$, we visualise two flow field snapshots (initial and final validation times) in figures 13 and 14, comparing the ROM prediction, the DNS and the projection of DNS onto the six POD modes. The ROM captures the flow structure well, including the separation bubble, although again, the wake exhibits higher errors due to phase shifts.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

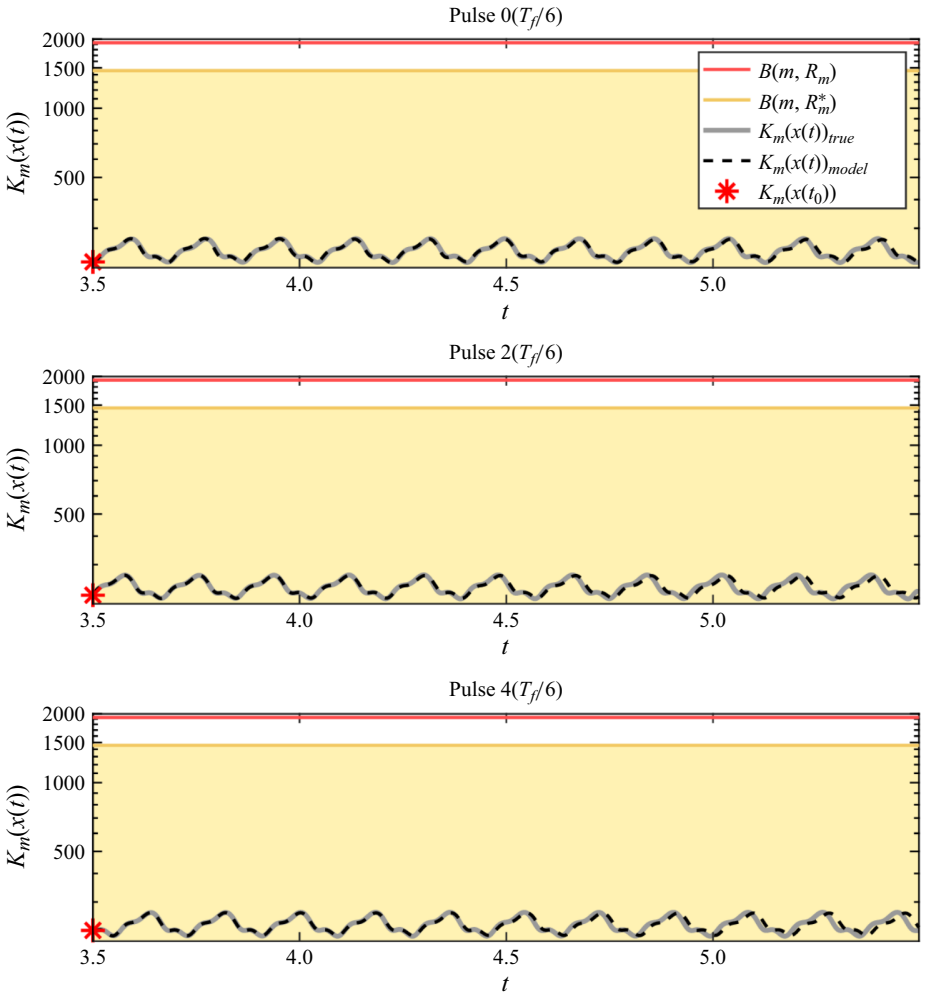


Figure 12. Validation of kinetic energy evolution for pulsed flow at $x/c = 0.3$ with actuation at $t = 0\,T_f/6$, $2\,T_f/6$ and $4\,T_f/6$. In each panel, the solid grey line shows the true energy, the black dashed line is the ROM prediction and the red star marks the initialisation point. The shaded yellow band indicates the trapping region $B(m, R_m^*)$ computed by our SDP method, while the red horizontal line denotes the larger trapping radius obtained via the Schlegel–Noack theorem.

Finally, we quantified the full-field reconstruction error across all test cases in figure 15. The ROM-predicted velocity field was compared with both the full DNS and the six-mode projection of DNS using the relative kinetic energy of the flow. Results confirm that the ROM closely approximates the separation bubble dynamics, although error accumulates downstream in the wake due to phase drift. This is consistent with prior observations and underscores the trade-off between model simplicity and full spatial accuracy in open flows Galletti *et al.* (2007), Callaham, Loiseau & Brunton (2023).

6. Conclusion

In this work we have presented a convex-optimisation-based framework for data-driven identification of ROMs of incompressible flows with certifiable boundedness properties. Our proposed trapping-SDP modelling framework (i) provides an unambiguous certificate

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

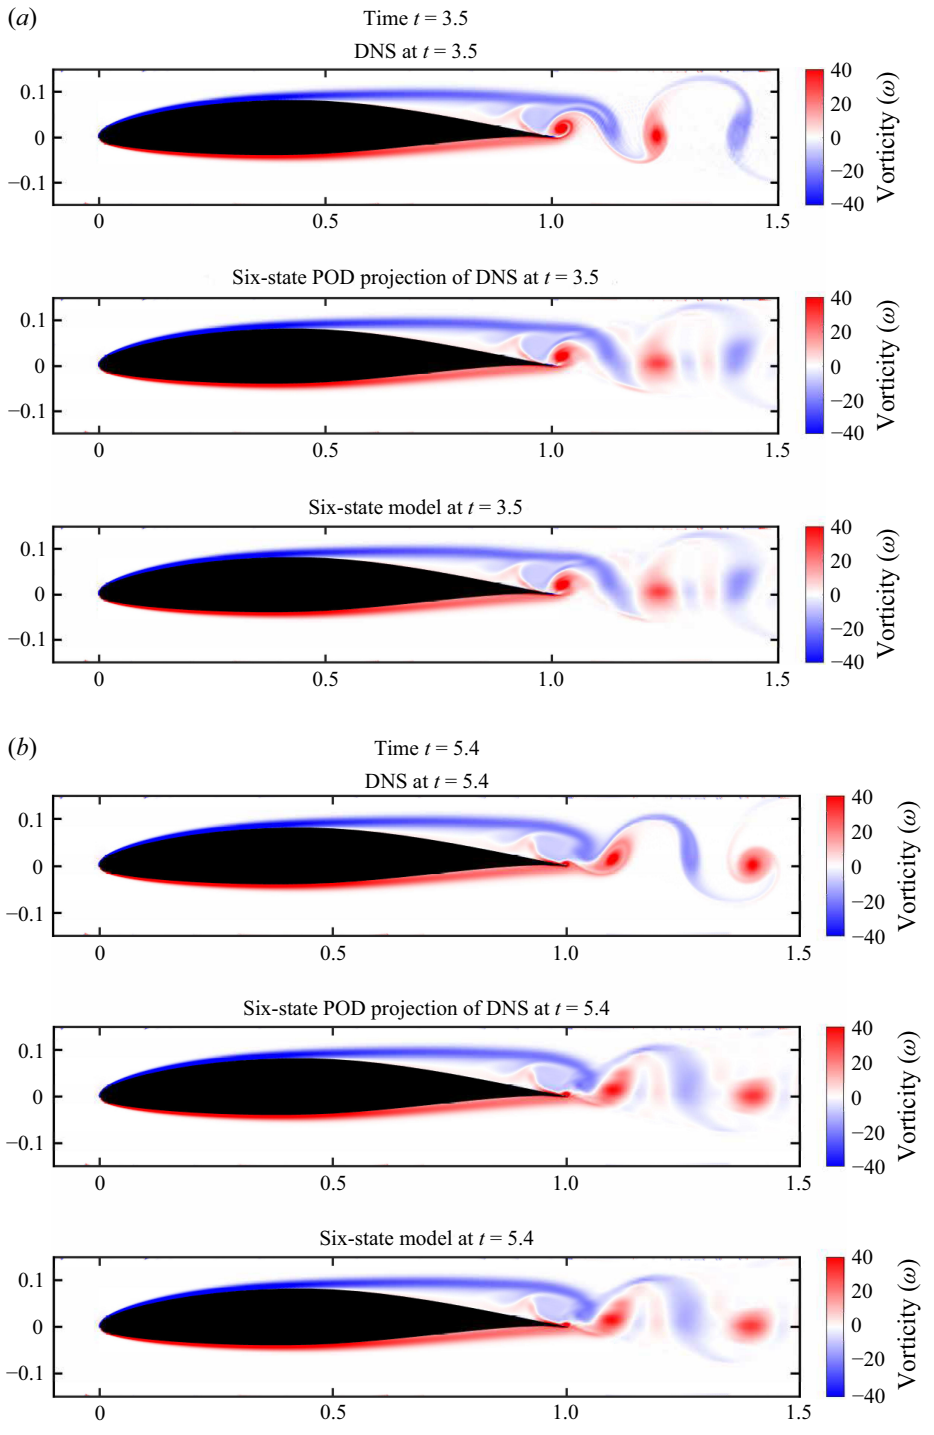


Figure 13. Vorticity fields for the pulsed case at $x/c = 0.3$ with actuation at $t = 0\ T_f/6$. Shown at (*a*) $t = 3.5$ and (*b*) $t = 5.4$ are (top) the full DNS solution, (middle) the six-state POD projection of the DNS and (bottom) the six-state ROM prediction. Results are shown for (*a*) $t = 3.5$ and (*b*) $t = 5.4$.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

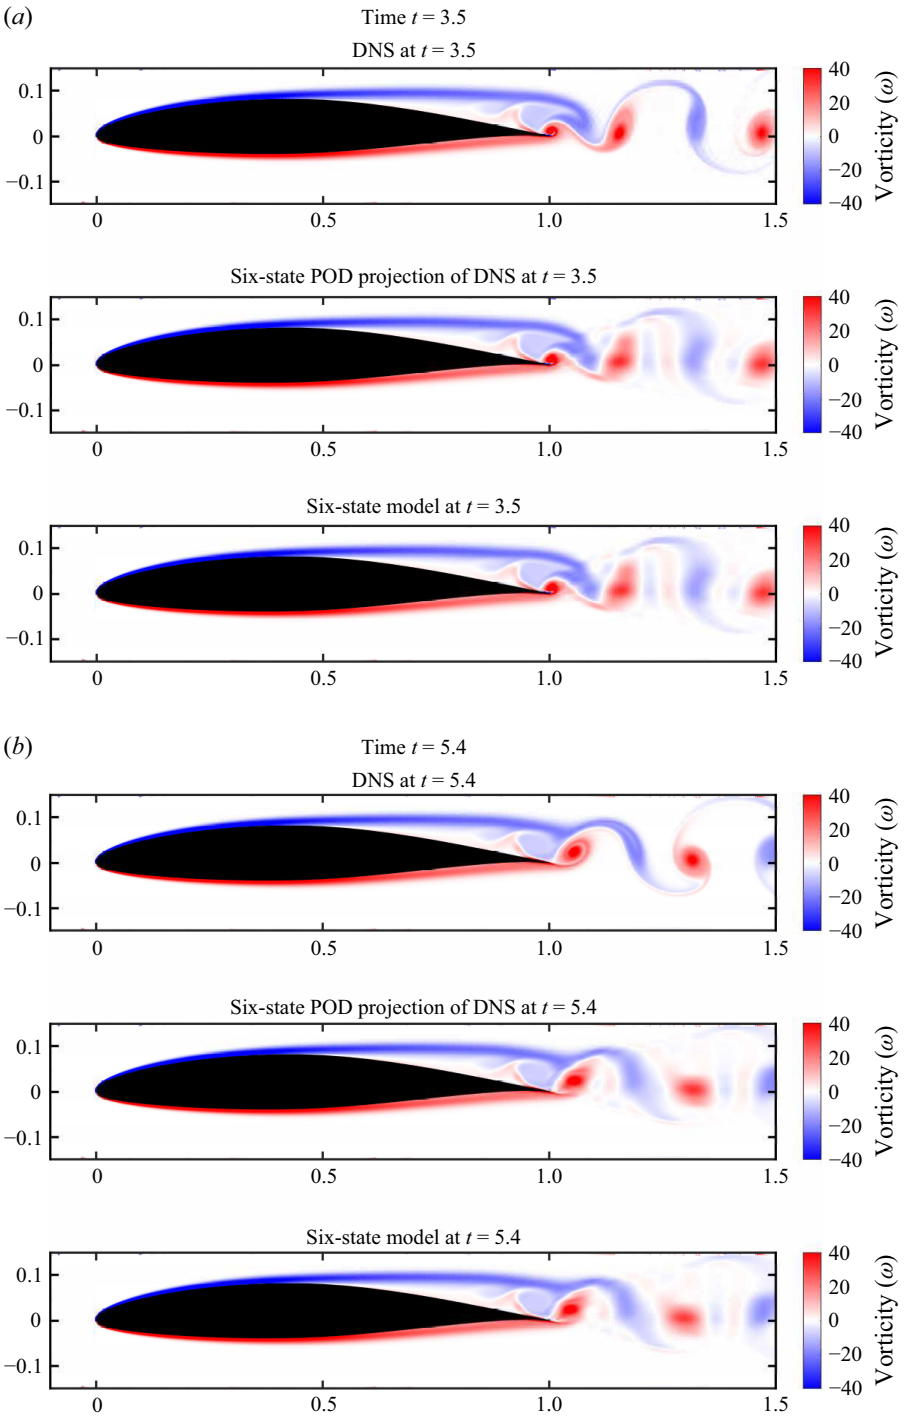


Figure 14. Vorticity fields for the pulsed case at $x/c = 0.3$ with actuation at $t = 4\,T_f/6$. Shown at (*a*) $t = 3.5$ and (*b*) $t = 5.4$ are (top) the full DNS solution, (middle) the six-state POD projection of the DNS and (bottom) the six-state ROM prediction. Results are shown for(*a*) $t = 3.5$ and (*b*) $t = 5.4$.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

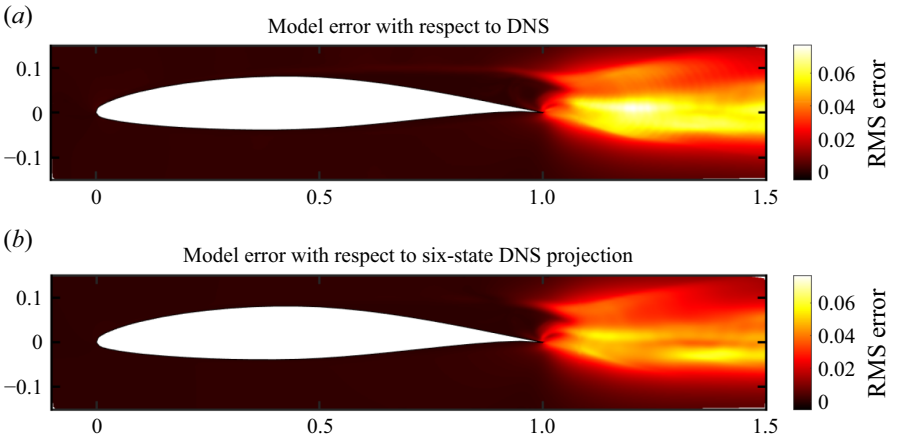


Figure 15. Time-averaged RMS error of the ROM over all validation cases. Panel (*a*) compares the ROM to full DNS and panel (*b*) to the six-state DNS projection. The error peaks in the wake downstream of the trailing edge, while the separation region above the aerofoil remains accurately captured.

of whether any lossless, globally bounded ROM exists for a given set of model parameters and hyperparameters, and (ii) if one does exist, computes a certified minimal Euclidean trapping ball $B(m, R_m^*)$ for the identified model within the adopted quadratic-energy certificate class (for the computed centre m). These convex subproblems are embedded into an alternating regression and stability-enforcement loop with provable convergence properties. We integrated our trapping-SDP loop within the context of the SINDy modelling framework to identify models that are simultaneously lossless, globally bounded and sparse. When the projected dynamics violate losslessness due to, for example, open-flow boundaries, numerical discretisation effects, truncation or weak compressibility, enforcing the lossless constraints in the regression should be interpreted as a modelling prior that regularises the surrogate model to yield bounded predictions. In such cases, the certificate applies rigorously to the learned ODE model rather than directly to the full-order system. Extensions to nonlinear-encoder coordinates or more general nonlinear representations would require an additional structure to recover analogous energy and losslessness properties. A remaining limitation is scalability: the present implementation uses off-the-shelf SDP solvers that enables modelling of systems with tens of modes. The size of the regression problem and the number of constraints grow with system dimension, introducing a computational bottleneck for higher-order systems. Such bottlenecks can potentially be overcome through tailored implementations that exploit problem structure, which is the subject of future investigation.

We demonstrated the effectiveness of our modelling approach first on two canonical benchmark problems: the chaotic Lorenz system and a nine-state Galerkin model of sinusoidal shear flow. In these examples, our approach recovered models that faithfully reproduced the known benchmark with regards to model coefficients and minimum trapping-region radius. We then applied our approach to obtain a six-mode model of unsteady separation over a NACA-65(1)-412 aerofoil at $Re = 20\,000$ and $\alpha = 4°$ from DNS data. The resulting model exhibited uniformly bounded long-time behaviour under arbitrary initial conditions, while retaining only the essential affine, linear and quadratic terms needed for accurate prediction over short- and long-time horizons.

The trapping-SDP modelling framework was shown to yield models with optimal trapping regions in all the examples considered. This distinguishing feature of the

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

trapping-SDP modelling framework facilitates the identification of models that are both physically and dynamically consistent with the governing equations, enabling predictive capabilities for both short- and long-time horizons. Moreover, by solving convex subproblems at each step, we were able to leverage mature off-the-shelf solvers to realise a tractable and reliable method for computing data-driven models with these certifiably optimal boundedness properties. We anticipate that the core idea – bringing convex, globally valid certificates into data-driven fluids modelling – will open new pathways toward trustworthy ROMs for prediction, optimisation and the control of complex flows.

Funding. This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-21-1-0434.

Declaration of interests. The authors report no conflict of interest.

Appendix A. Derivation of the lossless translated form

For completeness, we briefly derive (2.8). Starting from (2.7), we write

$$A(m) = L + 2\begin{bmatrix} m^\top Q^{(1)} \\ \vdots \\ m^\top Q^{(n)} \end{bmatrix} =: L + M(m). \tag{A1}$$

Taking the symmetric part gives

$$A_S(m) = \frac{1}{2}\big(A(m) + A(m)^\top\big) = L_S + \frac{1}{2}\big(M(m) + M(m)^\top\big). \tag{A2}$$

Now consider the $(i,\ j)$ entry of $M(m)$. Using the symmetry of $Q^{(i)}$, we obtain

$$M_{ij}(m) = 2\sum_{k=1}^{n} m_k Q^{(i)}_{kj} = 2\sum_{k=1}^{n} m_k Q^{(i)}_{jk}. \tag{A3}$$

Hence,

$$\frac{1}{2}\left(M_{ij}(m) + M_{ji}(m)\right) = \sum_{k=1}^{n} m_k \left(Q^{(i)}_{jk} + Q^{(j)}_{ik}\right). \tag{A4}$$

The lossless, or energy-preserving, condition on the quadratic tensor can be written as

$$Q^{(i)}_{jk} + Q^{(j)}_{ik} + Q^{(k)}_{ij} = 0, \tag{A5}$$

so that

$$Q^{(i)}_{jk} + Q^{(j)}_{ik} = -Q^{(k)}_{ij}. \tag{A6}$$

Substituting this identity into the previous expression yields

$$\frac{1}{2}\left(M_{ij}(m) + M_{ji}(m)\right) = -\sum_{k=1}^{n} m_k Q^{(k)}_{ij}. \tag{A7}$$

Therefore,

$$A_S(m) = L_S - \sum_{k=1}^{n} m_k Q^{(k)}, \tag{A8}$$

which is precisely (2.8). The factor of 2 in (2.7) is absorbed when taking the symmetric part, and the minus sign follows directly from the lossless identity above.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

Appendix B. Obtaining physics-based constraints

In order to obtain reduced-order governing equations that are consistent with the Navier–Stokes form, the physical properties of the incompressible fluid must be captured. The constrained approach described in § 2.4 allows active enforcement of desired conditions. To enforce physics-based constraints, the assumption that the quadratic and bilinear terms in the Navier–Stokes equation must be energy conserving is made. As the state x is chosen such that it is related directly to the kinetic energy of the perturbation, the constraint to make the quadratic nonlinear term energy preserving can be written in the form: $x^\top Q(x)x = 0$. We note that it is also possible to use the POD-basis representation defined in (5.3) together with the fact that POD modes form an orthonormal set to rewrite (2.1) in terms of the POD coefficients as

$$\dot{x}_i(t) = \sum_{j=1}^{n}\sum_{k=1}^{n} Q_{jk}^{(i)} x_j x_k + \sum_{j=1}^{n} L_{ij}x_j + c_i \qquad i = 1, 2, \ldots, n. \tag{B1}$$

The derivation of using the terms $Q_{jk}^{(i)}$, L_{ij} and c_i in (B1) to obtain constraints for Ξ is based on triadic interactions inherent to Galerkin projection methods (Rempfer & Fasel 1994; Balajewicz, Dowell & Noack 2013). The approach takes advantage of the symmetric nature of the term $Q_{jk}^{(i)}$ and the fact that it retains the energy-preserving property of the nonlinearity in the Navier–Stokes equations. The energy within each POD mode is expressed by $K_i(t) = (1/2)x_i^2(t)$, which is differentiated in time to obtain (Rempfer & Fasel 1994)

$$\dot{K}_i(t) = \sum_{j=1}^{n}\sum_{k=1}^{n} Q_{jk}^{(i)} x_i x_j x_k + \sum_{j=1}^{n} L_{ij}x_i x_j + c_i x_i \qquad i = 1, 2, \ldots, n. \tag{B2}$$

Here, the constraint on $Q_{jk}^{(i)}$ means that the quadratic nonlinearity can only serve to exchange energy between modes and does not contribute to the rate of change of the total energy of the system. Using this and the fact that the sum of the individual modes constitutes the total energy of the system, let $K(t) = \sum_{i=1}^{n} K_i(t) = (1/2)\sum_{i=1}^{n} x_i^2(t)$. The cubic contribution to $\dot{K}(t)$ is

$$0 = \sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{k=1}^{n} Q_{jk}^{(i)}\, x_i x_j x_k. \tag{B3}$$

Using the rules of index permutation and the symmetric property of $Q_{jk}^{(i)}$, we obtain the energy-preserving constraint

$$Q_{jk}^{(i)} + Q_{ki}^{(j)} + Q_{ij}^{(k)} = 0. \tag{B4}$$

Equation (B4) imposes the following three categories of constraints that are used to construct a constraint matrix C and d.

(i) Intrinsic constraint: occurs when $i = j = k$ and implies that $Q_{ii}^{(i)} = 0$ for any i. This is a result of a mode not being able to exchange energy with itself via the quadratic nonlinearity. This creates n constraints, and n rows are added to the matrix C.

(ii) Binary constraint: occurs when $i \neq j = k$ or $i = j \neq k$, and implies that a quadratic contribution involving x_j^2 in $\dot{x}_i$ must be balanced by a corresponding mixed term involving $x_i x_j$ in $\dot{x}_j$. Intuitively, this constraint arises when one mode interacts

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

quadratically with one other mode. To impose this, $n(n-1)$ rows are added to the matrix C (one for each possible binary constraint).

(iii) Extrinsic constraint: occurs when i, j and k are all distinct, and implies that a term involving $x_j x_k$ in $\dot{x}_i$ participates in a fully triadic interaction among modes i, j and k. This constraint describes energy exchange among three modes, where the exchange between two modes is mediated by a third. The complexity of these interactions causes the number of constraints added as rows to C to grow rapidly, i.e. $n!/3!(n-3)! = (1/6)n(n-1)(n-2)$.

A total of $n + n(n-1) + (1/6)n(n-1)(n-2)$ rows in C and elements in d are therefore added upon the first solution of the constraint matrices. Constraints for small entries in the coefficient matrix Ξ are then added on subsequent iterations until a solution is reached. This procedure results in constraints that ensure the quadratic nonlinearity conserves energy, consistent with the nonlinear physics of the incompressible Navier–Stokes equations.

Appendix C. Governing equations and numerical approximation

We consider the compressible Navier–Stokes equations for conservation of mass, momentum and energy, which can be written in non-dimensional form as a system of equations where the flux vector is divided into an advective (superscript a) and a viscous part (superscript v):

$$\frac{\partial \boldsymbol{Q}}{\partial t} + \boldsymbol{F}^a_x + \boldsymbol{G}^a_y - \frac{1}{Re_f}\left(\boldsymbol{F}^v_x + \boldsymbol{G}^v_y\right) = \boldsymbol{S}. \tag{C1}$$

Here the solution vector $\boldsymbol{Q} = [\,\rho \quad \rho u \quad \rho v \quad \rho E\,]^T$ and ρ, u, v and E denote the density and velocity in the x and y direction and the total energy, respectively. The equation of state closes the system: $p = \rho T/\gamma M^2$. All quantities are non-dimensionalised with respect to a characteristic, problem-dependent length scale, reference velocity, density and temperature yielding the non-dimensional Reynolds number, Re, and Mach number, M. At a sufficiently low Mach number, the flow can be modelled as incompressible to good approximation – a fact that we exploit to obtain bounded low-order models within our data-driven modelling framework.

The system of equations is spatially approximated using a discontinuous Galerkin spectral element method and integrated in time with a fourth-order explicit Runge–Kutta scheme. Gauss–Lobatto quadrature nodes are used for the spatial integration and a kinetic energy conserving split-form approximation of the advective volume fluxes ensures stability of the scheme through cancellation of aliasing errors from the nonlinear terms. For a detailed description of the scheme, we refer the reader to Kopriva (2009), Gassner, Winters & Kopriva (2016) and Klose, Jacobs & Kopriva (2019).

C.1. *The DNS of pulsed flow over an aerofoil*

The flow over a NACA 65(1)-412 aerofoil is simulated at a Reynolds number based on the chord length of $Re = 20\,000$ and a Mach number of $M = 0.3$. The low Mach number ensures a nominally incompressible flow.

The computational domain is given in figure 16 and consists of 2256 quadrilateral elements. The domain dimensions prevent blockage and follow the recommendations in Nelson, Jacobs & Kopriva (2016). The wall-boundary elements are curved and fitted to a spline representing the aerofoil's surface according to Nelson *et al.* (2016). The boundary conditions are specified weakly on the fluxes according to a Riemann solver (Jacobs,

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

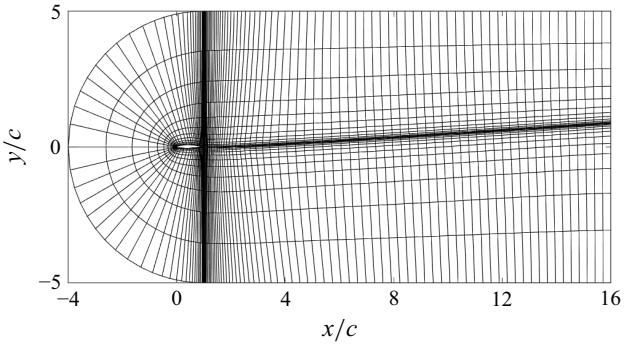


Figure 16. Two-dimensional computational domain for NACA 65(1)-412. Only elements without interior Gauss–Lobatto nodes are shown.

Kopriva & Mashayek 2002). The solution vector is approximated with a twelfth-order polynomial, giving a total of 381 264 collocation points in the domain.

The compressible Navier Stokes equations can be written as

$$\frac{\partial \boldsymbol{U}}{\partial t} + \nabla \cdot \boldsymbol{F} = \boldsymbol{S}, \tag{C2}$$

where the solution vector $\boldsymbol{U}$ is

$$\boldsymbol{U} = [\rho \;\; \rho u \;\; \rho v \;\; \rho e]^T, \tag{C3}$$

where the first entry accounts for mass conservation, the second and third entries for momentum conservation on x and y, respectively, and the fourth entry for the total energy conservation. The flux tensor is $\boldsymbol{F}$ (see Nelson *et al.* 2016 for more details on the compressible formulation).

Here we include a source term, which for a pulse, $\boldsymbol{S}$, is specified following Suzuki, Colonius & Pirozzoli (2004), Kamphuis *et al.* (2018), taking the form

$$\boldsymbol{S}(t, \boldsymbol{x}) = \boldsymbol{b}(t) \odot \frac{\boldsymbol{a}}{\sqrt{(2\pi)^n |\boldsymbol{\Sigma}^{-1}|}} \exp\left[-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{c})^T \boldsymbol{\Sigma}^T (\boldsymbol{x} - \boldsymbol{c})\right], \tag{C4}$$

where $\odot$ denotes the Hadamard product, $\boldsymbol{\Sigma}$ is a 2×2 covariance matrix, $\boldsymbol{a}$ a vector of amplitudes and $\boldsymbol{c}$ identifies the centre coordinates of the Gaussian distribution. We set $\boldsymbol{\Sigma}$ to be

$$\boldsymbol{\Sigma} = \begin{matrix} \dfrac{\cos^2\theta}{2\sigma_x^2} + \dfrac{\sin^2\theta}{2\sigma_y^2} & \dfrac{\sin 2\theta}{4\sigma_x^2} - \dfrac{\sin 2\theta}{4\sigma_y^2} \\ \dfrac{\sin 2\theta}{4\sigma_x^2} - \dfrac{\sin 2\theta}{4\sigma_y^2} & \dfrac{\sin^2\theta}{2\sigma_x^2} + \dfrac{\cos^2\theta}{2\sigma_y^2} \end{matrix} \tag{C5}$$

for $\theta \in [-45°, 45°]$, so that the principal axes of the Gaussian function are rotated by an angle θ with respect to the y axis, and with variances given by

$$\sigma_x = 0.0012/\sqrt{2}; \qquad \sigma_y = 0.0096/\sqrt{2}. \tag{C6}$$

The vector of amplitudes, $\boldsymbol{a}$, is set as

$$\boldsymbol{a} = \left[0, \;\; \sqrt{(2\pi)^n \left|\Sigma_{mj}^{-1}\right|}, \;\; \sqrt{(2\pi)^n \left|\Sigma_{mj}^{-1}\right|}, \;\; \sqrt{(2\pi)^n \left|\Sigma_{mj}^{-1}\right|}\right], \tag{C7}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

x/c	0.1	0.2	0.3	0.4	0.45	0.5
T	$0(T_f/6)$	$1(T_f/6)$	$2(T_f/6)$	$3(T_f/6)$	$4(T_f/6)$	$5(T_f/6)$

Table 2. Locations, x/c, on the suction side of the aerofoil for the source actuation and respective times, T, as fractions of the base flow natural period T_f.

and the time-dependent component, $\boldsymbol{b}(t)$, as

$$\boldsymbol{b}(t) = \delta(t-T)[0, \ -\sin(\theta), \ \cos(\theta), \ u(x,t)b_2(t) + v(x,t)b_3(t)\ ]^T, \tag{C8}$$

where $\delta(\tau)$ is the Dirac delta function. Note that the source is zero mass and the angle θ was chosen so that the principal axis of the Gaussian function is perpendicular to the aerofoil surface. The Dirac delta conveniently integrates in time to the Heaviside function, which means that numerically the pulsed case is implemented exactly by superposing a Gaussian distribution onto the solution vector at time T.

The data required as input for the analysis of the optimal location and time for a pulsed control is obtained by conducting DNS and systematically sampling the parameter space spanned by $\boldsymbol{c}$ and T. Table 2 collates the values for which simulations are conducted. Here, the source centre $\boldsymbol{c}$ in (C4) is located at a distance of $2\sigma_y$ into the flow, measured perpendicular to the aerofoil suction side, at the listed cord-length locations (following Suzuki *et al.* 2004), and times are presented as sixth's of the natural period for the lift. Combinations of all locations and activation times for the source give us a total of 36 test points.

Appendix D. Coefficient data for the six-mode model of unsteady separation

This appendix reports coefficients for the affine, linear and quadratic terms associated with the six-mode model of unsteady separation over a NACA 65(1)-412 aerofoil at $Re = 20\,000$ (with respect to the chord) and an angle of attack of $\alpha = 4°$ identified and discussed in § 5.1. Note the reported values are not specified to full floating point precision here. The trapping-region size for the model coefficients reported below is $R_m^* = 1{,}452$. The modeled coefficients are

$$c = \begin{bmatrix} 141.8727 & 0 & 0 & 0 & 145.0715 & 229.5386 \end{bmatrix}^\top, \tag{D1}$$

$$L = \begin{bmatrix} 1.1316 & 28.7305 & 2.1558 & 0 & 0 & -1.338 \\ -27.3208 & -0.23667 & 0 & 0 & -2.717 & 0 \\ -2.5646 & 0 & 0 & 0 & -30.6392 & 0 \\ -7.2196 & 0 & 6.5418 & 0 & 0 & -26.1932 \\ -2.8727 & 0 & 0 & -5.107 & -17.9816 & -94.1002 \\ 0 & 3.9537 & 0 & 0 & 103.1758 & -12.6275 \end{bmatrix}, \tag{D2}$$

$$Q^{(1)} = \begin{bmatrix} 0 & -0.011056 & -0.17693 & -0.1263 & 0.094271 & -0.019469 \\ -0.011056 & -0.077894 & -0.16887 & 0.27678 & -0.05589 & -0.17187 \\ -0.17693 & -0.16887 & -0.51254 & 0.031921 & -0.17917 & 0.0041533 \\ -0.1263 & 0.27678 & 0.031921 & -0.44694 & -0.066242 & -0.075209 \\ 0.094271 & -0.05589 & -0.17917 & -0.066242 & 0.057274 & 0.010509 \\ -0.019469 & -0.17187 & 0.0041533 & -0.075209 & 0.010509 & 0.21571 \end{bmatrix}, \tag{D3}$$

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

$$Q^{(2)} = \begin{bmatrix} 0.022111 & 0.038947 & -0.058039 & 0.12261 & -0.12096 & 0.051903 \\ 0.038947 & 0 & 0.18141 & 0.11448 & -0.031486 & 0.12183 \\ -0.058039 & 0.18141 & 0 & 0.087642 & 0.077644 & -0.055879 \\ 0.12261 & 0.11448 & 0.087642 & -0.1374 & 0.080832 & 0.16085 \\ -0.12096 & -0.031486 & 0.077644 & 0.080832 & 0.0055095 & -0.12572 \\ 0.051903 & 0.12183 & -0.055879 & 0.16085 & -0.12572 & -0.015538 \end{bmatrix}, \tag{D4}$$

$$Q^{(3)} = \begin{bmatrix} 0.35387 & 0.22691 & 0.25627 & -0.055247 & 0.016723 & 0 \\ 0.22691 & -0.36282 & 0 & 0.13667 & -0.065643 & 0.052092 \\ 0.25627 & 0 & 0 & 0 & 0 & 0.18596 \\ -0.055247 & 0.13667 & 0 & 0 & -0.28978 & -0.040504 \\ 0.016723 & -0.065643 & 0 & -0.28978 & -0.058656 & 0 \\ 0 & 0.052092 & 0.18596 & -0.040504 & 0 & 0.011922 \end{bmatrix}, \tag{D5}$$

$$Q^{(4)} = \begin{bmatrix} 0.2526 & -0.39939 & 0.023326 & 0.22347 & 0.031266 & 0.01276 \\ -0.39939 & -0.22896 & -0.22431 & 0.068698 & -0.073167 & -0.10128 \\ 0.023326 & -0.22431 & 0 & 0 & -0.0021628 & -0.058946 \\ 0.22347 & 0.068698 & 0 & 0 & 0.077367 & -0.047753 \\ 0.031266 & -0.073167 & -0.0021628 & 0.077367 & -0.42712 & 0.038952 \\ 0.01276 & -0.10128 & -0.058946 & -0.047753 & 0.038952 & -0.54495 \end{bmatrix}, \tag{D6}$$

$$Q^{(5)} = \begin{bmatrix} -0.18854 & 0.17685 & 0.16245 & 0.034976 & -0.028637 & 0.11008 \\ 0.17685 & 0.062971 & -0.012001 & -0.0076647 & -0.0027547 & 0.15125 \\ 0.16245 & -0.012001 & 0 & 0.29194 & 0.029328 & -0.042638 \\ 0.034976 & -0.0076647 & 0.29194 & -0.15473 & 0.21356 & 0 \\ -0.028637 & -0.0027547 & 0.029328 & 0.21356 & 0 & -0.041723 \\ 0.11008 & 0.15125 & -0.042638 & 0 & -0.041723 & 0.087508 \end{bmatrix}, \tag{D7}$$

$$Q^{(6)} = \begin{bmatrix} 0.038939 & 0.11996 & -0.0041533 & 0.062449 & -0.12059 & -0.10785 \\ 0.11996 & -0.24367 & 0.0037873 & -0.059568 & -0.025528 & 0.0077689 \\ -0.0041533 & 0.0037873 & -0.37192 & 0.09945 & 0.042638 & -0.0059612 \\ 0.062449 & -0.059568 & 0.09945 & 0.095507 & -0.038952 & 0.27248 \\ -0.12059 & -0.025528 & 0.042638 & -0.038952 & 0.083446 & -0.043754 \\ -0.10785 & 0.0077689 & -0.0059612 & 0.27248 & -0.043754 & 0 \end{bmatrix}. \tag{D8}$$

The shift vector m for the reported model is

$$m = \begin{bmatrix} -126.9980 & 9.8762 & 37.1703 & -78.8471 & -141.3316 & -118.1777 \end{bmatrix}^{\top}. \tag{D9}$$

References

Ahmadi, A.A. & El Khadir, B. 2023 Learning dynamical systems with side information. *SIAM Rev.* **65** (1), 183–223.

Balajewicz, M.J., Dowell, E.H. & Noack, B.R. 2013 Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier–Stokes equation. *J. Fluid Mech.* **729**, 285–308.

Benner, P., Gugercin, S. & Willcox, K. 2015 A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* **57** (4), 483–531.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

BOYD, S., EL GHAOUI, L., FERON, E. & BALAKRISHNAN, V. 1994 *Linear Matrix Inequalities in System and Control Theory*. SIAM.
BOYD, S. & VANDENBERGHE, L. 2004 *Convex Optimization*. Cambridge University Press.
BRUNTON, S.L. & NOACK, B.R. 2015 Closed-loop turbulence control: progress and challenges. *Appl. Mech. Rev.* **67** (5), 050801.
BRUNTON, S.L., NOACK, B.R. & KOUMOUTSAKOS, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.
BRUNTON, S.L., PROCTOR, J.L., KUTZ, J.N. & BIALEK, W. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci. USA* **113**, 3932–3937.
BUZHARDT, J. & GRAHAM, M.D. 2025 Identifying efficient routes to laminarization: an optimization approach. arXiv:2508.08519.
CALLAHAM, J.L., LOISEAU, J.-C. & BRUNTON, S.L. 2023 Multiscale model reduction for incompressible flows. *J. Fluid Mech.* **973**, A3.
CAVALIERI, A.V.G. & NOGUEIRA, P.A.S. 2022 Reduced-order Galerkin models of plane Couette flow. *Phys. Rev. Fluids* **7**, L102601.
CAZEMIER, W., VERSTAPPEN, R. & VELDMAN, A.E.P. 1998 Proper orthogonal decomposition and low-dimensional models for the Burgers' equation. *J. Comput. Phys.* **146** (2), 451–473.
CHAMPION, K., ZHENG, P., ARAVKIN, A.Y., BRUNTON, S.L. & KUTZ, J.N. 2020 A unified sparse optimization framework to learn parsimonious physics-informed models from data. *IEEE Access* **8**, 169259–169271.
CHERNYSHENKO, S.I., GOULART, P.J., HUANG, D. & PAPACHRISTODOULOU, A. 2014 Polynomial sum of squares in fluid dynamics: a review with a look ahead. *Phil. Trans. R. Soc. Lond. A: Math. Phys. Engng Sci.* **372** (2020), 20130350.
DOERING, C.R. & GIBBON, J.D. 1995 *Applied Analysis of the Navier-Stokes Equations*. Cambridge Texts in Applied Mathematics. Cambridge University Press.
DURAISAMY, K., IACCARINO, G. & XIAO, H. 2019 Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**, 357–377.
ENAYATI KAFSHGARKOLAEI, A. & HEMATI, M.S. 2025 Local stability and stabilization of quadratic-bilinear systems using Petersen's lemma. pre-print arXiv:2503.21040.
GALLETTI, B., BOTTARO, A., BRUNEAU, C.-H. & IOLLO, A. 2007 Accurate model reduction of transient and forced wakes. *Eur. J. Mech. B/Fluids* **26** (3), 354–366.
GASSNER, G.J., WINTERS, A.R. & KOPRIVA, D.A. 2016 Split form nodal discontinuous Galerkin schemes with summation-by-parts property for the compressible Euler equations. *J. Comput. Phys.* **327**, 39–66.
GOLUSKIN, D. 2020 Bounding extrema over global attractors using polynomial optimisation. *Nonlinearity* **33** (9), 4878–4899.
GOULART, P.J. & CHERNYSHENKO, S. 2012 Global stability analysis of fluid flows using sum-of-squares. *Phys. D: Nonlinear Phenom.* **241** (6), 692–704.
GOYAL, P., DUFF, I.P. & BENNER, P. 2024 Guaranteed stable quadratic models and their applications in sindy and operator inference. pre-print arXiv:2308.13819v2.
HEIDE, A.L. & HEMATI, M.S. 2025 Optimization framework for analyzing nonlinear stability due to sparse finite-amplitude perturbations. *Phys. Rev. Fluids* **10**, 014401.
HOLMES, P., LUMLEY, J.L., BERKOOZ, G. & ROWLEY, C.W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.
HUANG, D., CHERNYSHENKO, S., GOULART, P.J., LASAGNA, D., TUTTY, O. & FUENTES, F. 2015 Sum-of-squares of polynomials approach to nonlinear stability of fluid flows: an example of application. *Proc. R. Soc. Lond. A: Math. Phys. Engng Sci.* **471** (2183), 20150622.
JACOBS, G., KOPRIVA, D. & MASHAYEK, F. 2002 Outflow boundary conditions for the multidomain staggered-grid spectral method for the Navier–Stokes equations. In *40th AIAA Aerospace Sciences Meeting & Exhibit*. *AIAA Paper* 2002-0903. AIAA.
KALUR, A., MUSHTAQ, T., SEILER, P. & HEMATI, M.S. 2022 Estimating regions of attraction for transitional flows using quadratic constraints. *IEEE Control Syst. Lett.* **6**, 482–487.
KAMPHUIS, M., JACOBS, G.B., CHEN, K., SPEDDING, G. & HOEIJMAKERS, H. 2018 Pulse actuation and its effects on separated Lagrangian coherent structures for flow over a cambered airfoil. In *2018 AIAA Aerospace Sciences Meeting*. *AIAA Paper* 2018-2255. AIAA.
KAPTANOGLU, A.A., CALLAHAM, J.L., ARAVKIN, A., HANSEN, C.J. & BRUNTON, S.L. 2021 Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Phys. Rev. Fluids* **6** (9), 094401.
KAPTANOGLU, A.A. *et al.* 2022 PySINDy: a comprehensive Python package for robust sparse system identification. *J. Open Source Softw.* **7** (69), 3994.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

KARNIADAKIS, G.E., KEVREKIDIS, I.G., LU, L., PERDIKARIS, P., WANG, S. & YANG, L. 2021 Physics-informed machine learning. *Nat. Rev. Phys.* **3** (6), 422–440.
KHALIL, H.K. 2002 *Nonlinear Systems*, 3rd edn. Prentice Hall.
KLOSE, B.F., JACOBS, G.B. & KOPRIVA, D.A. 2019 On the robustness and accuracy of marginally resolved discontinuous Galerkin schemes for two dimensional Navier–Stokes flows. In *AIAA Scitech 2019 Forum*.
KOPRIVA, D.A. 2009 *Implementing Spectral Methods for Partial Differential Equations*. Springer.
KRAMER, B. 2021 Stability domains for quadratic-bilinear reduced-order models. *SIAM J. Appl. Dyn. Syst.* **20** (2), 981–996.
KRAMER, B., PEHERSTORFER, B. & WILLCOX, K.E. 2024 Learning nonlinear reduced models from data with operator inference. *Annu. Rev. Fluid Mech.* **56**, 521–548.
LASAGNA, D., HUANG, D., TUTTY, O.R. & CHERNYSHENKO, S. 2016 Sum-of-squares approach to feedback control of laminar wake flows. *J. Fluid Mech.* **809**, 628–663.
LIAO, S.-C., LEONID HEIDE, A., HEMATI, M.S. & SEILER, P.J. 2025 A convex optimization approach to compute trapping regions for lossless quadratic systems. *Intl J. Robust Nonlinear Control* **35** (6), 2425–2436.
LIU, C. & GAYME, D.F. 2020 Input–output inspired method for permissible perturbation amplitude of transitional wall-bounded shear flows. *Phys. Rev. E* **102** (6), 063108.
LOISEAU, J.-C. & BRUNTON, S.L. 2018 Constrained sparse Galerkin regression. *J. Fluid Mech.* **838**, 42–67.
LORENZ, E.N. 1963 Deterministic nonperiodic flow. *J. Atmos. Sci.* **20** (2), 130–141.
LU, L., JIN, P., PANG, G., ZHANG, Z. & KARNIADAKIS, G.E. 2021 Learning operators with deeponet. arXiv:1910.03193.
MOEHLIS, J., FAISST, H. & ECKHARDT, B. 2004 A low-dimensional model for turbulent shear flows. *New J. Phys.* **6** (1), 56.
NELSON, D.A., JACOBS, G.B. & KOPRIVA, D.A. 2016 Effect of boundary representation on viscous, separated flows in a discontinuous-Galerkin Navier–Stokes solver. *Theor. Comput. Fluid Dyn.* **30**, 363–385.
PEHERSTORFER, B. & WILLCOX, K. 2016 Data-driven operator inference for nonintrusive projection-based model reduction. *Comput. Meth. Appl. Mech. Engng* **306**, 196–215.
PENG, M., MANOHAR, K., BRUNTON, S., KAPTANOGLU, A. & HANSEN, C. 2024 Local stability guarantees for data-driven quadratically nonlinear models. arXiv:2403.00324v2.
RAISSI, M., PERDIKARIS, P. & KARNIADAKIS, G.E. 2019 Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707.
REMPFER, D. & FASEL, H.F. 1994 Dynamics of three-dimensional coherent structures in a flat-plate boundary layer. *J. Fluid Mech.* **275**, 257–283.
ROWLEY, C.W. & DAWSON, S.T.M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.
SCHLEGEL, M. & NOACK, B.R. 2015 On long-term boundedness of Galerkin models. *J. Fluid Mech.* **765**, 325–352.
SCHLOSSER, C. & KORDA, M. 2021 Converging outer approximations to global attractors using semidefinite programming. *Automatica* **134**, 109900.
SCHMIDT, O.T. 2020 Bispectral mode decomposition of nonlinear flows. *Nonlinear Dyn.* **102**, 1–23.
DE SILVA, B., CHAMPION, K., QUADE, M., LOISEAU, J.-C., KUTZ, J. & BRUNTON, S. 2020 PySINDy: a Python package for the sparse identification of nonlinear dynamical systems from data. *J. Open Source Softw.* **5** (49), 2104.
SUZUKI, T., COLONIUS, T. & PIROZZOLI, S. 2004 Vortex shedding in a two-dimensional diffuser: theory and simulation of separation control by periodic mass injection. *J. Fluid Mech.* **520**, 187–213.
TAIRA, K., BRUNTON, S.L., DAWSON, S.T.M., ROWLEY, C.W., COLONIUS, T., MCKEON, B.J., SCHMIDT, O.T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L.S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4411.
TAIRA, K., HEMATI, M.S., BRUNTON, S.L., SUN, Y., DURAISAMY, K., BAGHERI, S., DAWSON, S.T.M. & YEH, C.-A. 2020 Modal analysis of fluid flows: applications and outlook. *AIAA J.* **58** (3), 998–1022.
TANK, J.D., KLOSE, B.F., JACOBS, G.B. & SPEDDING, G.R. 2021 Flow transitions on a cambered airfoil at moderate Reynolds number. *Phys. Fluids* **33** (9), 093105.
THUEREY, N., WEISSENOW, K., PRANTL, L. & HU, W. 2020 Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows. *AIAA J.* **58** (1), 25–36.
TSENG, P. 2001 Convergence of a block coordinate descent method for nondifferentiable minimization. *J. Optim. Theor. Applics* **109** (3), 475–494.

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press

https://doi.org/10.1017/jfm.2026.11671 Published online by Cambridge University Press