



A priori sparsification of Galerkin models

Riccardo Rubini¹, Davide Lasagna^{1,†} and Andrea Da Ronch¹

¹Faculty of Engineering and Physical Sciences, University of Southampton, SO17 1BJ Southampton, UK

(Received 10 June 2021; revised 7 March 2022; accepted 27 March 2022)

A methodology to generate sparse Galerkin models of chaotic/unsteady fluid flows containing a minimal number of active triadic interactions is proposed. The key idea is to find an appropriate set of basis functions for the projection representing elementary flow structures that interact minimally one with the other, resulting in a triadic interaction coefficient tensor with sparse structure. Interpretable and computationally efficient Galerkin models are obtained, since a reduced number of triadic interactions are computed to evaluate the right-hand side of the model. To find the basis functions, a subspace rotation technique is used, whereby a set of proper orthogonal decomposition (POD) modes is rotated into a POD subspace of larger dimension using coordinates associated with low-energy dissipative scales to alter energy paths and the structure of the triadic interaction coefficient tensor. This rotation is obtained as the solution of a non-convex optimisation problem that maximises the energy captured by the new basis, promotes sparsity and ensures long-term temporal stability of the sparse Galerkin system. We demonstrate the approach on two-dimensional lid-driven cavity flow at $Re = 2 \times 10^4$ where the motion is chaotic and energy interactions are scattered in modal space. We show that the procedure generates Galerkin models with a reduced set of active triadic interactions, distributed in modal space according to established knowledge of scale interactions in two-dimensional flows. This property, however, is observed only if long-term temporal stability is included explicitly in the formulation, indicating that a dynamical constraint is necessary to obtain a physically consistent sparsification.

Key words: computational methods, low-dimensional models, machine learning

1. Introduction

From a physical standpoint, turbulence is a multi-scale phenomenon exhibiting a wide hierarchy of spatial and temporal scales. This property, coupled with the intrinsic nonlinearity of the governing equations, poses considerable difficulties to its modelling and analysis. One of the major challenges to obtain a satisfactory mathematical description

† Email address for correspondence: davide.lasagna@soton.ac.uk

of this phenomenon arises from the fact that the dynamics of flow structures at a particular length or time scale cannot be examined in isolation without also considering at the same time the whole hierarchy of complementary scales (Domaradzki *et al.* 1994). In fact, nonlinear interactions between triads of scales play a fundamental role as they are the main driver of energy transfer between coherent structures (Pope 2001; Moffatt 2014). In turn, the organisation of triadic interactions has a direct influence on the physics of a number of flow phenomena, such as direct and inverse energy cascades (Kolmogorov 1991), or transition to turbulence (Craik 1971; Rempfer & Fasel 1994*a,b*) in different flow configurations (Schmidt 2020). Overall, this property makes the development of computationally efficient and physically interpretable reduced-order dynamical models a challenging task.

Historically, the study of triadic interactions has been conducted by first employing an appropriate decomposition technique to educe coherent structures from the turbulent motion and then characterising the intensity of the intermodal couplings through the perspective of the resulting Galerkin model (Noack *et al.* 2008). For homogeneous isotropic turbulence, Fourier modes provide an optimal representation (Brasseur & Wei 1994; Laval, Dubrulle & Nazarenko 1999), but for flows in complex geometries, modes identified from data with proper orthogonal decomposition (POD) (Rempfer & Fasel 1994*a*; Couplet, Sagaut & Basdevant 2003) or with frequency domain decomposition methods (Schmid 2010; Towne, Schmidt & Colonius 2018; Symon, Illingworth & Marusic 2020) have been often adopted. One of the key findings of such studies is that energy transfers are not uniformly distributed in modal space. In fact, not all interactions have the same importance and energy flows along preferential directions. Specifically, there is evidence suggesting that the nonlinear interaction pattern among coherent structures is often sparse. In other words, the evolution of structures at a certain length scale depends predominantly upon a subset of all other structures (Kraichnan 1971; Ohkitani 1990; Brasseur & Wei 1994), and the influence of interactions with the complementary set of structures can be neglected generally with minor global effects. This behaviour has been observed for a multiplicity of different flows, ranging from bluff body wakes (Jin, Symon & Illingworth 2021) to transitional boundary layers (Rempfer & Fasel 1994*a,b*) and separated flows (Couplet *et al.* 2003). However, these studies have focused typically on the interaction between modes educed using a predetermined flow decomposition technique that does not necessarily capture faithfully or optimally the aforementioned physics of scale interactions, since no information regarding nonlinear mechanisms that may produce such interactions is utilised in the decomposition. Hence rationalising scale interactions in these models and identifying relevant physical mechanisms may be challenging, especially in high-dimensional systems Schmidt (2020).

In a previous work (Rubini, Lasagna & Da Ronch 2020*b*), we utilised l_1 -based regression methods (Brunton, Proctor & Kutz 2016; Loiseau & Brunton 2018) to extract sparsity patterns in the intermodal energy transfers in large Galerkin models of multi-scale flows, to construct sparse, computationally efficient and interpretable models.

This technique consists of calibrating all model coefficients by solving a regression problem (Cordier, El Majd & Favier 2010) augmented with a term that penalises the l_1 norm of the coefficients vector. This type of penalisation is known to promote sparse solutions (Tibshirani 1996, 2013), and thus results in calibrated models where weak triadic interactions that do not contribute significantly to the overall dynamics are pruned. However, a rigorous connection between the modal structures and the projection model is necessarily lost, since the procedure involves calibrating model coefficients without modifying the basis functions. Hence the analysis of energy paths in the model is not

necessarily consistent with the spatial structure of the basis functions. In addition, it was demonstrated that the organisation and sparsity of energy interactions is not invariant with respect to a change of the basis functions utilised for the analysis of the energy budget.

One avenue in this direction that has been explored recently is to utilise sparse coding methods (Olshausen & Field 1996) and sparse principal component analysis (Jolliffe, Trendafilov & Uddin 2003) to encode sparsity directly in the spatio-temporal structure of the reduced basis. For instance, Deshmukh *et al.* (2016) utilised sparse coding techniques to develop a modified version of POD, producing a compact representation that spans all scales of the observed data. To this end, the authors augmented the POD variational statement with a sparsity-promoting l_1 penalisation term on the temporal coefficients matrix. This forces the temporal coefficients to display a marked intermittent character to capture local dynamics, i.e. periods of activity alternated with periods of low amplitude, with only a subset of the entire set active at any instant. However, the orthogonality between modes is not necessarily preserved. Erichson *et al.* (2020) generalised the same approach and proposed more efficient computational algorithms for its solution that include orthogonality constraints for the spatial modes. The authors demonstrated that this formulation can be a useful diagnostic tool for problems with rich spatio-temporal behaviour, since it can provide a cleaner data description by identifying localised spatial structures in the data and can disambiguate between distinct time scales. However, spatially or temporally localised modal structures obtained using these techniques may not result directly in a sparser triadic interaction pattern once Galerkin projection is performed. In fact, these techniques utilise only flow data, and it is unclear if this is sufficient to capture the sparse nature of scale interactions.

In this paper, we attempt to bridge this gap. We propose a novel sparsification method in which the goal is to seek directly modal structures that capture energy transfer mechanisms efficiently. More precisely, we seek a new basis set that produces a sparse quadratic coefficient tensor regulating triadic interactions, without the need for an *a posteriori* tuning based on l_1 -based regression. The computational approach is inspired by the subspace rotation technique of Balajewicz, Dowell & Noack (2013) and Balajewicz, Tezaur & Dowell (2016), where a small rotation of the POD subspace was sought to absorb the unresolved dissipative scales into the basis functions and stabilise the long-term behaviour without the need for empirical eddy-viscosity terms. Here, the key idea is to seek a small rotation of the original POD subspace to alter and sparsify energy transfer paths. In practice, the rotation is found by solving a constrained optimisation problem, minimising the loss of energy optimality subject to a constraint on the sparsity of the quadratic coefficient tensor. We refer to this method as *a priori* sparsification, since sparse characteristics are obtained before the projection step, not after.

As a demonstration, we use two-dimensional incompressible lid-driven cavity flow with uniform lid velocity at Reynolds number $Re = 2 \times 10^4$, where fluid motion is chaotic (Auteri, Parolini & Quartapelle 2002). As opposed to flows at lower Reynolds number just beyond bifurcation to time-dependent flow, energy transfers between modal structures obtained with POD are scattered in modal space and dense Galerkin systems with full model coefficient tensors are obtained. This dynamically complex flow exhibiting a ‘light turbulent’ regime represents a suitable benchmark problem to understand how the proposed optimisation procedure captures the sparsity inherent to the energy cascade of a multi-scale flow. The two-dimensional lid-driven cavity flow is also an established test case for the development and validation of model order reduction techniques (Cazemier, Verstappen & Veldman 1998; Terragni, Valero & Vega 2011; Balajewicz *et al.* 2013; Deshmukh *et al.* 2016; Arbabi & Mezić 2017; Fick *et al.* 2018), and we thus consider it here as an exemplar to demonstrate our approach.

The paper is organised as follows. Section 2 outlines the general methodology to generate reduced-order models by Galerkin projection and how energy interactions between modes defining such models can be analysed. Subsequently, the subspace rotation technique to generate sparse models is described. Results are reported in § 3.

2. Methodology

2.1. Galerkin models and energy analysis

We consider the space of square integrable solenoidal velocity vector fields defined over a spatial domain Ω , endowed by the standard inner product

$$(\mathbf{u}, \mathbf{v}) := \int_{\Omega} \mathbf{u} \cdot \mathbf{v} \, d\Omega, \quad (2.1)$$

where \mathbf{u}, \mathbf{v} are two elements of such space. The resulting $\mathcal{L}^2(\Omega)$ norm is denoted as $\|\mathbf{u}\|_2 = \sqrt{(\mathbf{u}, \mathbf{u})}$. Using the time-averaged velocity field $\bar{\mathbf{u}}(\mathbf{x})$ as a base flow, and denoting by $\mathbf{u}'(t, \mathbf{x}) = \mathbf{u}(t, \mathbf{x}) - \bar{\mathbf{u}}(\mathbf{x})$ the velocity fluctuation, the N -dimensional Galerkin ansatz

$$\mathbf{u}(t, \mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}) + \mathbf{u}'(t, \mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}) + \sum_{i=1}^N a_i(t) \boldsymbol{\phi}_i(\mathbf{x}) \quad (2.2)$$

is introduced to describe the space–time velocity field, where $a_i(t)$ and $\boldsymbol{\phi}_i(\mathbf{x})$, $i = 1, \dots, N$, are the temporal and global spatial modes, respectively. We assume throughout that the spatial modes form an orthonormal set and satisfy homogeneous boundary conditions on Ω . Reduced-order models are then derived by projecting the Navier–Stokes equations for incompressible flows onto the subspace defined by the spatial modes (Rowley & Dawson 2017) using the inner product defined in (2.1).

Considering configurations where the boundaries are either no-slip walls or periodic and do not change over time, the pressure term arising from the projection vanishes identically for solenoidal modes (Schlegel & Noack 2015; Lee & Dowell 2020). The system of coupled nonlinear ordinary differential equations (ODEs)

$$\dot{a}_i(t) = C_i + \sum_{j=1}^N L_{ij} a_j(t) + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk} a_j(t) a_k(t), \quad i = 1, \dots, N, \quad (2.3)$$

is then obtained, defining the temporal evolution of the coefficients $a_i(t)$. Here, the tensors $\mathbf{C} \in \text{Re}^N$, $\mathbf{L} \in \text{Re}^{N \times N}$ and $\mathbf{Q} \in \text{Re}^{N \times N \times N}$ are defined by suitable inner products involving the spatial modes. In particular, the coefficients are

$$C_i = \frac{1}{Re} (\bar{\mathbf{u}}, \nabla^2 \bar{\mathbf{u}}) - (\bar{\mathbf{u}}, \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}), \quad (2.4a)$$

$$L_{ij} = \frac{1}{Re} (\boldsymbol{\phi}_i, \nabla^2 \boldsymbol{\phi}_j) - (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j \cdot \nabla \bar{\mathbf{u}}) - (\boldsymbol{\phi}_i, \bar{\mathbf{u}} \cdot \nabla \boldsymbol{\phi}_j), \quad (2.4b)$$

$$Q_{ijk} = -(\boldsymbol{\phi}_i, \boldsymbol{\phi}_j \cdot \nabla \boldsymbol{\phi}_k). \quad (2.4c)$$

The expansion (2.2) provides a suitable foundation to examine interactions between coherent structures in complex geometries. Here, we follow established approaches (Rempfer & Fasel 1994b) and analyse such interactions by introducing the modal energies $E_i(t) = \frac{1}{2} a_i(t) a_i(t)$, $i = 1, \dots, N$. For an expansion consisting of N modes, the domain

integral of the kinetic energy of velocity fluctuation expressed by the ansatz (2.2) is then given by

$$E(t) = \sum_{i=1}^N E_i(t). \tag{2.5}$$

The instantaneous rate of change of the modal energies is given by

$$\dot{E}_i(t) = C_i a_i(t) + \sum_{j=1}^N L_{ij} a_i(t) a_j(t) + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk} a_i(t) a_j(t) a_k(t), \quad i = 1, \dots, N, \tag{2.6}$$

obtained by multiplying (2.3) by $a_i(t)$. The right-hand side of (2.6) is composed of three terms describing energy transfers between the hierarchy of modes (Noack *et al.* 2008; Noack, Morzynski & Tadmor 2011). The first two describe variations of energy due to production/dissipation arising from interactions with the mean flow and from viscous effects. The third term defines variations of energy arising from inviscid nonlinear interactions between triads of modes. Additional insight can be gained by taking the temporal average of (2.6), i.e. by examining the time-averaged energy budget of system (2.3). Exploiting that the temporal coefficients have zero mean as POD modes have vanishing mean, we obtain

$$\sum_j L_{ij} \overline{a_i a_j} + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk} \overline{a_i a_j a_k} = 0, \quad i = 1, \dots, N, \tag{2.7}$$

where the summation over j in the linear term accounts for the case where temporal coefficients are not uncorrelated in time. Equation (2.7) shows that, on average, the energy balance is regulated by the production/dissipation described in the linear term and by the nonlinear energy transfer rate described by the quadratic term. As explained in Balajewicz *et al.* (2013), the residual of (2.7) vanishes only when $N \rightarrow \infty$ and it is generally expected to be positive for finite-dimensional Galerkin models.

Following Rempfer & Fasel (1994a), to better visualise the relative importance of the triadic interactions in an average, we introduce the tensor $\mathbf{N} \in \text{Re}^{N \times N \times N}$ with entries

$$N_{ijk} = Q_{ijk} \overline{a_i a_j a_k}, \tag{2.8}$$

defining the the average nonlinear transfer rate between triads of modes. The study of the organisation and structure of this tensor and how these are altered by the sparsification algorithm is the main focus of this work.

2.2. Subspace rotation technique

To identify the new set of modal structures, we utilise a subspace rotation technique introduced in the context of stabilisation of Galerkin models by Balajewicz *et al.* (2013) (see also Amsallem & Farhat 2012). Geometrically, this technique consists of seeking a rotation of an N -dimensional POD subspace within a larger POD subspace of dimension M . The rotation is defined by a transformation matrix $\mathbf{X} \in \text{Re}^{M \times N}$, satisfying $\mathbf{X}^T \mathbf{X} = \mathbf{I}$ to ensure that the rotated spatial basis functions remain an orthonormal set. The rotated basis functions and the associated temporal coefficients, denoted in what follows with a tilde,

are expressed as a linear combination of the original POD spatial and temporal modes as

$$\tilde{\phi}_i(\mathbf{x}) = \sum_{j=1}^M X_{ji} \phi_j(\mathbf{x}), \tag{2.9a}$$

$$\tilde{a}_i(t) = \sum_{j=1}^M X_{ji} a_j(t). \tag{2.9b}$$

It is worth pointing out that finding a new set of modal structures directly would be a much higher dimensional problem to tackle. The number of unknowns would be proportional to the numbers of modes sought multiplied by the number of degrees of freedom of the problem at hand. Seeking new modal structures as a linear combination of POD modes represents a significant reduction in complexity, controllable by varying the dimension M . Using POD modes as building blocks also has the advantage of producing a basis with good energy reconstruction properties.

The linear and quadratic coefficients of the Galerkin system (2.3) obtained by projection on the rotated subspace are then given by the matrix expressions

$$\tilde{\mathbf{C}} = \mathbf{X}^T \mathbf{C} \quad \text{and} \quad \tilde{\mathbf{L}} = \mathbf{X}^T \mathbf{L} \mathbf{X}, \tag{2.10a,b}$$

while the quadratic coefficients are cubic polynomial functions of the entries of the rotation matrix

$$\tilde{Q}_{ijk} = \sum_{p,q,r=1}^M Q_{pqr} X_{pi} X_{qj} X_{rk}, \quad i, j, k = 1, \dots, N. \tag{2.11}$$

In these expressions, the tensors $\mathbf{C} \in \text{Re}^M$, $\mathbf{L} \in \text{Re}^{M \times M}$ and $\mathbf{Q} \in \text{Re}^{M \times M \times M}$ are the Galerkin coefficient tensors obtained from the M -dimensional set of original POD modes.

Our goal is to seek a rotation matrix for which the rotated quadratic interaction coefficient tensor $\tilde{\mathbf{Q}}$ has a sparse structure, i.e. where as many as possible of the quadratic interaction coefficients are identically zero.

At this stage, it is worth noting that any rotation is necessarily accompanied by a loss of average fluctuation kinetic energy reconstructed by the new basis. The energy reconstructed by a set of P POD modes can be quantified by utilising the average modal energies $\lambda_i = \overline{a_i a_i}$, $i = 1, \dots, P$. Arranging them into a diagonal matrix $\mathbf{A}_P \in \text{Re}^{P \times P}$, the trace $\text{Tr}(\mathbf{A}_P)$ defines an upper bound for the reconstructed energy for any P -dimensional set of basis functions, due to well-known optimality properties of POD. Similarly, the energy reconstructed by the rotated basis can be expressed with the average modal energies $\tilde{\lambda}_i = \overline{\tilde{a}_i \tilde{a}_i}$, $i = 1, \dots, N$, of the rotated temporal coefficients (2.9b) and arranging them into the diagonal matrix $\tilde{\mathbf{A}}_N = \mathbf{X}^T \mathbf{A}_M \mathbf{X} \in \text{Re}^{N \times N}$. The loss of reconstructed average fluctuation kinetic energy with respect to an N -dimensional POD subspace is then quantified as

$$\mathcal{J}(\mathbf{X}) = \text{Tr}(\mathbf{A}_N - \mathbf{X}^T \mathbf{A}_M \mathbf{X}). \tag{2.12}$$

As observed in Balajewicz *et al.* (2013), this quantity is necessarily non-negative due to the optimality of the original POD basis, i.e.

$$\overline{\int_{\Omega} \|\mathbf{u}'(t, \mathbf{x})\|^2 d\Omega} = \text{Tr}(\mathbf{A}_{\infty}) > \text{Tr}(\mathbf{A}_N) \geq \text{Tr}(\tilde{\mathbf{A}}_N), \tag{2.13}$$

where the last equality holds for $M > N$ in trivial cases only. In addition, the quantity (2.12) is also always identically zero when $M = N$, for any \mathbf{X} , since any linear combination of N POD modes necessarily spans the same original N -dimensional subspace.

To measure sparsity of the quadratic interaction coefficient tensor, we use the l_1 norm operator, denoted as $\|\cdot\|_1$, in light of the practical intractability of the zero norm in optimisation (Jovanović, Schmid & Nichols 2014). Then the trade-off between energy optimality and sparsity is expressed by formulating the following constrained optimisation problem:

$$\min_{\mathbf{X}} \quad \text{Tr}(\Lambda_N - \mathbf{X}^T \Lambda_M \mathbf{X}), \tag{2.14a}$$

$$\text{subject to} \quad \|\tilde{\mathbf{Q}}\|_1 \leq \|\mathbf{Q}\|_1 / \xi, \tag{2.14b}$$

$$\mathbf{X}^T \mathbf{X} = \mathbf{I}_{N \times N}, \tag{2.14c}$$

referred to as problem P1 in what follows. The role of the objective function (2.14a) is to favour transformation matrices that minimise the loss of energy optimality, producing a set of basis functions with good energy representation ability, as in Balajewicz *et al.* (2016). On the other hand, the constraint (2.14b) encourages sparse solutions, because some entries of $\tilde{\mathbf{Q}}$ are shrunk to zero during the solution of (2.14) by the non-differentiability of the l_1 norm (Friedman, Hastie & Tibshirani 2008; Tibshirani 2013). Sparsification of the linear coefficients tensor $\tilde{\mathbf{L}}$ should not be expected. Then the weight ξ is an arbitrary penalisation parameter that controls the relative l_1 norm of the rotated quadratic coefficient tensor and the sparsity of the resulting Galerkin model. This parameter ranges from 1 to arbitrarily large values. When it is equal to 1, the l_1 norm of the rotated tensor is not affected and no sparsification is obtained. Conversely, using larger ξ constrains the norm of the rotated tensor to decrease. This promotes sparsity due to the non-differentiability of the l_1 norm, with a mechanism (see Appendix A) similar to well-known LASSO methods (Tibshirani 1996). In addition, different definitions of the sparsification constraint (2.14b) can be used to highlight different physical aspects of the resulting bases. As an example, in an exploratory study we observed that applying the l_1 norm constraint on $\tilde{\mathbf{N}}$ instead of $\tilde{\mathbf{Q}}$, it is possible to promote sparsity in the average triadic interaction tensor. Interestingly, with this formulation, sparsification is promoted by penalising the entries $\tilde{a}_i \tilde{a}_j \tilde{a}_k$ more than the entries of $\tilde{\mathbf{Q}}$. As a result, the temporal coefficients display intermittency behaviour similar to what is observed by Deshmukh *et al.* (2016).

It is worth noting that using the *a posteriori* LASSO-based sparsification method (Brunton *et al.* 2016; Rubini *et al.* 2020b), all quadratic coefficients could, in principle, be set to zero by using a large regularisation weight in the LASSO optimisation problem. This is because the constant, linear and quadratic coefficients of the Galerkin model are directly the optimisation variables. In the present case, the model coefficients cannot be modified directly, but the model tuning is performed indirectly through the rotation matrix \mathbf{X} , which is the actual optimisation variable of the problem. The important consequence is that it might not always be possible to find a rotation that sets an arbitrarily large number of model coefficients to zero. This indicates that problem (2.14) might not have a feasible solution if the penalisation weight is too large.

An important characteristic of optimisation problem (2.14) is that while the objective (2.14a) is convex, the sparsity-promoting constraint (2.14b) is not, as it involves cubic polynomials in the optimisation variables, the entries of the transformation matrix \mathbf{X} . Consequently, the solution might not be unique, and several local minima, corresponding to different sets of basis functions, may be obtained by starting the optimisation from different initial guesses. However, as demonstrated in Appendix B, starting the

optimisation from small random perturbations of the original POD basis (justified by the need to retain good energy reconstruction properties) produced consistently the same optimal solutions, which will be presented in § 3.

Assuming that a feasible solution of problem (2.14) can be found, the Galerkin model constructed from projection onto the optimal rotated basis does not necessarily possess better long-term temporal stability characteristics than the original POD model. In fact, it is well known that POD-Galerkin models exhibit long-term instability because of the deficit of energy dissipation attributed to the truncation of small dissipative scales (Noack *et al.* (2008), Schlegel & Noack (2015); see also the recent work of Grimberg, Farhat & Youkilis (2020) for a point against this argument). In the present case, the transformation \mathbf{X} obtained from solution of (2.14) does not necessarily result in an improved description of dissipative processes. Classically, this issue is cured by introducing, *a posteriori*, an eddy-viscosity-type term in the Galerkin model (Galletti *et al.* 2004; Noack, Papas & Monkewitz 2005; Östh *et al.* 2014). However, an *a posteriori* correction would not remain in the spirit of the present work. We thus favour the subspace-rotation-based stabilisation approach proposed by Balajewicz *et al.* (2013, 2016), which can be introduced naturally in the present formulation. In practice, we augment problem (2.14) with the additional implicit constraint

$$\text{Tr}(\tilde{\mathbf{L}}) = \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = -\eta, \tag{2.15}$$

where the auxiliary variable $\eta \in \text{Re}^+$ is chosen such that

$$\chi(\eta) = \frac{\overline{E(t)} - \overline{E_{DNS}(t)}}{\overline{E_{DNS}(t)}} = 0, \tag{2.16}$$

i.e. such that the relative difference of the average fluctuation kinetic energies from direct numerical simulation (DNS) and from numerical simulation of the new model vanishes. The variable η controls dissipation mechanisms in the Galerkin model by altering the spectrum of $\tilde{\mathbf{L}}$ and ensures long-term stability. As observed by Balajewicz *et al.* (2013), η is not known *a priori*, but can be found in an inner optimisation loop to ensure that the excess average fluctuation kinetic energy defined by $\chi(\eta)$ is zero. With this additional constraint, problem (2.14) becomes

$$\min_{\mathbf{X}} \text{Tr}(\mathbf{A}_N - \mathbf{X}^T \mathbf{A}_M \mathbf{X}), \tag{2.17a}$$

$$\text{subject to } \|\tilde{\mathbf{Q}}\|_1 \leq \|\mathbf{Q}\|_1 / \xi, \tag{2.17b}$$

$$\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}) = -\eta \quad \text{with } \chi(\eta) = 0, \tag{2.17c}$$

$$\mathbf{X}^T \mathbf{X} = \mathbf{I}_{N \times N}. \tag{2.17d}$$

In this formulation, denoted as P2 henceforth, there is still only one free parameter, ξ . The additional constraint guarantees long-term stability but it can be satisfied (i.e. the problem is feasible) only if $M > N$. In fact, any rotation $\mathbf{X} \in \text{Re}^{N \times N}$ cannot alter the subspace spanned by the original N POD modes and the associated description of dissipation mechanisms captured by the model. In what follows, we consider models with ratio $M/N = 2$ and 3.

In practice, a small random perturbation of the original N -dimensional POD basis was used as initial guess for the optimisation, and successive optimisation problems for different penalisation ξ were started from the solution of the previous problem. Problem (2.17) was solved with the open-source package for nonlinear and non-convex optimisation NLOpt (Johnson 2014). We utilised a solver implementing the method of

moving asymptotes (MMA) algorithm (Svanberg 2014), which requires the gradient of the objective function and of the constraints. The key element to make the procedure viable is to evaluate the sparsity-promoting constraint and its gradient as efficiently as possible. A naive implementation requires $O(M^3N^3)$ operations for the evaluation of the sparsity-promoting constraint, and $O(M^4N^4)$ for the evaluation of its gradient with respect to the rotation \mathbf{X} , and costs become intractable quickly. A significantly more efficient algorithm to compute these two quantities – with costs scaling as $O(MN^3 + M^2N^2 + M^3N)$ for the evaluation of the sparsity-promoting constraint and its gradient – can be derived and is key to making the procedure viable (see discussion of the method in Appendix C).

One further technical remark is that the sparsity-promoting constraint is a non-smooth function of the transformation matrix \mathbf{X} , posing difficulties for the utilisation of gradient-based optimisation algorithms. One approach is to implement a subgradient descent method, often used for the solution of the LASSO method (Friedman *et al.* 2008). However, this algorithm can be implemented only if an analytical solution of the optimisation problem is known, which is not the case here. In this work, we used a manual soft-thresholding approach where entries of the rotated tensor $\tilde{\mathbf{Q}}$ smaller than the numerical tolerance specified to the gradient-based optimiser (typically $tol = 10^{-5}$) are set to zero at the end of the optimisation.

The Galerkin models obtained from solution of (2.17) are then characterised by examining the density of the rotated triadic interaction tensor, defined as

$$\rho = \frac{\|\tilde{\mathbf{Q}}\|_0}{\|\mathbf{Q}\|_0}, \tag{2.18}$$

where the l_0 norm $\|\cdot\|_0$ counts the non-zero elements of a tensor. The density can also be expressed as the average

$$\rho = \frac{1}{N} \sum_{i=1}^N \rho_i, \tag{2.19}$$

with the modal densities $\rho_i = \|\tilde{\mathbf{Q}}_i\|_0 / \|\mathbf{Q}_i\|_0$, $i = 1, \dots, N$, being the relative number of non-zero coefficients in the slices \mathbf{Q}_i of the quadratic coefficient tensor associated with each modal index. To express the energy captured by the N -dimensional rotated basis, we also introduce the global energy reconstruction factor

$$e_N = \text{Tr}(\tilde{\mathbf{A}}_N) / \text{Tr}(\mathbf{A}_\infty), \tag{2.20}$$

a quantity always strictly lower than 1. The density and the reconstruction factor of the rotated Galerkin models depend on the dimensions M and N , and on the penalisation weight ξ . To characterise the effects of these parameters, we visualise the rotated systems on the ρ – e_N plane, similar to the approach used in Rubini *et al.* (2020b).

3. Demonstration: two-dimensional lid-driven unsteady cavity flow

We now apply this methodology to two-dimensional unsteady flow in a lid-driven square cavity. This is the same test case that we utilised in our previous work (Rubini *et al.* 2020b) to demonstrate the properties of l_1 -based sparsification of Galerkin models.

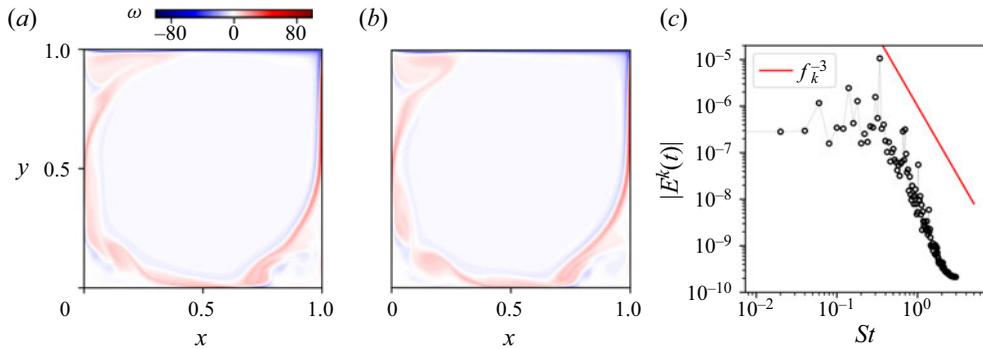


Figure 1. (a,b) Instantaneous vorticity fields for two different flow snapshots. (c) Amplitude of the Fourier transform of the fluctuating kinetic energy signal (2.5).

3.1. Problem definition and proper orthogonal decomposition

The Reynolds number is defined as $Re = LU/\nu$, where L and U are the cavity dimension and the (uniform) lid velocity, respectively, and ν is the kinematic viscosity. All physical variables introduced later are scaled with L , U and combinations thereof. We consider the flow regime establishing at $Re = 2 \times 10^4$, where the motion is chaotic (Auteri *et al.* 2002; Peng, Shiau & Hwang 2003). The domain is defined by the non-dimensional Cartesian coordinates $\mathbf{x} = (x, y)$, and the velocity vector $\mathbf{u}(t, \mathbf{x})$ is defined by the components $u(t, \mathbf{x})$ and $v(t, \mathbf{x})$. For visualisation purposes, we introduce the out-of-plane vorticity component $\omega = \partial v/\partial x - \partial u/\partial y$.

Numerical simulations were performed in OpenFOAM with the unsteady incompressible flow solver icofoam. The convective and viscous terms are discretised spatially with a second-order finite-volume technique, and the temporal term with a semi-implicit Crank–Nicolson scheme. A grid independence study was performed, examining average and unsteady flow quantities on increasingly finer meshes. The final mesh is composed of 300×300 cells, with refinement at the four cavity boundaries. This mesh is sufficiently fine to resolve the unsteady high-shear regions bounding the main vortex and the high vorticity filaments characteristic of two-dimensional turbulence, as well as the spatial structure of the lowest energy POD modes utilised in this study. Similar grid resolutions have been used by Cazemier *et al.* (1998) at similar Reynolds numbers.

Two snapshots of the vorticity field obtained from these simulations are shown in figures 1(a,b). Most of the dynamically interesting features in this regime originate at the bottom-right corner of the cavity. Specifically, the secondary vortex in the recirculation zone is shed erratically, producing wave-like disturbances advected along the shear layer bounding the primary vortex. These structures produce a strong quasi-periodic oscillation of the turbulent kinetic energy, as shown in figure 1(c), where the amplitude of the Fourier transform of the fluctuation kinetic energy signal (2.5) is shown, as a function of the Strouhal number $St = fL/U$, with f the dimensional frequency. Due to the chaotic nature of the flow, the energy spectrum has a strong broadband component with energy uniformly distributed across all the frequencies/spatial structures of the flow. This implies that unlike for periodic flows (Noack *et al.* 2011; Symon *et al.* 2020), the mean triadic interaction tensor \mathbf{N} is dense, with energy transfers highly scattered in modal space. However, the wave-like motion characterising the shear layer dynamics produces a strong energy peak at a characteristic non-dimensional Strouhal number $St = 0.7$ and its harmonics.

N	1	5	10	20	30	40	60	80	90	120
e_N	0.26	0.74	0.9	0.96	0.97	0.98	0.985	0.99	0.995	0.998

Table 1. Normalised cumulative energy reconstruction of the original POD basis functions.

From these simulations, we extract $N_T = 1000$ snapshots after initial transients have decayed, using a non-dimensional sampling period $\Delta t = 0.1$. These settings are sufficient to time-resolve adequately the fastest scales of motion as well as to include many shedding events originating at the bottom right corner. The snapshot POD of this dataset is then performed (Sirovich 1987). The normalised cumulative reconstructed energy (2.20) is reported in table 1 as a function of the mode number N .

3.2. Analysis of energy interaction of the original POD-Galerkin models

Before examining the properties of the sparsified Galerkin models, we first present some aspects of energy transfers of Galerkin models constructed using the original POD modes, using a reference model size $N = 30$. This dimension is chosen as a compromise between the need to observe a sufficient complexity of modal interactions and the need to limit computational costs associated with the solution of the optimisation problem (see scaling of costs in Appendix C). The POD model for $N = 30$ resolves approximately 97 % of the average fluctuation kinetic energy, which is sufficient to observe a wide dynamic range in the amplitude of the temporal coefficients and in the strength of energy interactions.

Figure 2(a) shows the organisation of the quadratic coefficients \mathbf{Q} . The slice for mode $i = 1$ is used as an illustrative example; other modes have similar characteristics. The coefficient tensor is dense as most coefficients are non-zero. No clear structure in the distribution of the coefficients can be observed, except for a slight asymmetry where coefficients for $k > j$ are often larger in magnitude, for mode $i = 1$ but also for most of the other modes. This asymmetry is a consequence of the definition of the quadratic coefficient tensor (2.4c), characterising convective transport mechanisms, and the range of spatial length scales described by the POD modes, as already observed in Rubini *et al.* (2020b). One slice of the average energy transfer rate tensor \mathbf{N} is reported in figure 2(b) using the base 10 logarithm of the magnitude. All entries of \mathbf{N} are generally non-zero, although the intermodal transfers in the Galerkin model are highly organised and the intensity of interactions varies across several orders of magnitude. This is the combined result of the structure of the quadratic coefficients tensor and the complex spectral structure of the temporal coefficients.

To further characterise energy paths, we split the interactions into four regions, denoted as LL, LH, HL and HH, with L and H denoting low and high index modes, respectively. These four regions represent a coarse-grained grouping of energy transfers between scales resolved by the model, assuming that low-index modes map to the largest scales of motion in the cavity, and high-index modes describe small-scale, low-energy features. This property has been verified for certain problems (Couplet *et al.* 2003; see also discussion in Grimberg *et al.* 2020) and is confirmed in the present case by visual inspection of the POD modes. We then compute the sum of the magnitudes of the average transfer rates contained in these four regions for each mode, to characterise in a coarse-grained fashion energy transfers between triads of high-energy modes and smaller, dissipative scales.

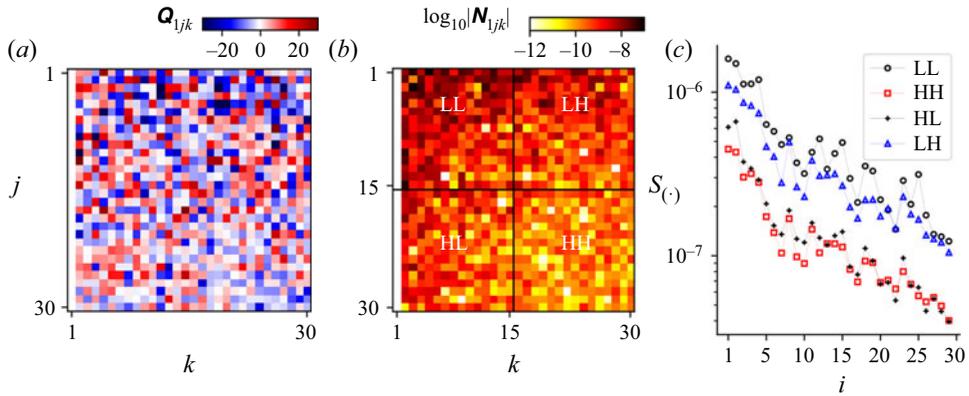


Figure 2. Organisation of (a) the quadratic coefficients \mathbf{Q} , and (b) the average nonlinear transfer rate \mathbf{N} , for mode $i = 1$. (c) Collective energy transfer rates associated with the four regions defined in (b), as functions of the modal index i .

For instance, for the region LL we compute

$$S_{LL} = \sum_{j=1}^n \sum_{k=1}^n |N_{ijk}|. \tag{3.1}$$

The subscript on S in expression (3.1) denotes one of the four regions. The other regions are defined by analogous expressions, differing in the range spanned by the indices j and k . Here, we arbitrarily select the cut-off at half of the spectrum ($n = 15$), but other choices are possible and do not change the following results qualitatively. For graphical convenience, we define the notation $S_{(\cdot)}$ where the dot represents the quantity (3.1) for the four regions defined in figure 2(b). The result of this analysis is shown in figure 2(c), as a function of the modal index i . First, energy transfers in the regions LL, HL and LH are generally more intense than those in region HH. This follows from the observation that large-scale/large-scale and large-scale/small-scale interactions are more relevant with respect to the small-scale/small-scale interactions, across the entire hierarchy and in a mean sense. In addition, we observe that the interactions LH are always more intense than the HL interactions. This is a consequence of the asymmetry observed previously in the tensor \mathbf{Q} , and not of the temporal coefficients, due to definition (2.8), where the indices j and k commute. This lack of symmetry is in agreement with the picture of energy transfers between scales in homogeneous isotropic two-dimensional turbulence (Ohkitani 1990; Laval *et al.* 1999), where the large scales interact with the small ones in a non-local fashion.

3.3. Model sparsification

We now consider models with dimension $N = 30$, for ratios $M/N = 2$ and 3, and examine in more detail the effect of the sparsity-promoting constraint (2.17b). For each ratio, a family of models with different density and reconstructed average kinetic energy is generated by increasing the penalisation parameter ξ . Optimal solutions are displayed on the $1/\xi - \|\tilde{\mathbf{Q}}\|_1 / \|\mathbf{Q}\|_1$ plane in figure 3(a). The red dashed line separates solutions that satisfy the sparsity-promoting constraint (2.17b) (white area feasibility region) from solutions that do not (red area). Note that the stability constraint is satisfied for all points

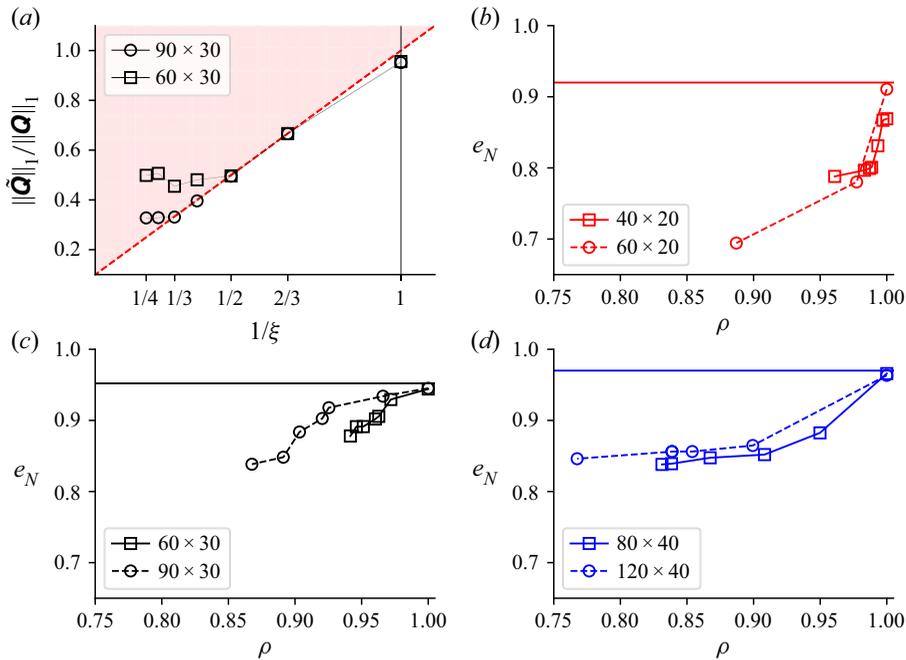


Figure 3. (a) Visualisation of the sparsity-promoting constraint (2.17b) on the $1/\xi - \|\tilde{\mathbf{Q}}\|_1 / \|\mathbf{Q}\|_1$ plane. The red region denotes the infeasibility set. (b–d) Plots of ρ versus e_N for families of models with dimensions $N = 30, 20, 40$ and ratios $M/N = 2, 3$ (open squares and open circles, respectively). The labels in the legends indicate values of $M \times N$.

reported in this figure, as we have noted that the optimiser is still able to satisfy (2.17c) when it first fails to satisfy the sparsity-promoting constraint (2.17b).

For small penalisation weights, the optimisation problem has feasible solutions that fall on the boundary of the feasibility region. This suggests that rotations of the original POD basis that minimise the energy loss are found on the boundary, i.e. there is competition between sparsification and energy representation. The key feature of figure 3(a) is that there exists a threshold value ξ_t above which the optimisation problem terminates unsuccessfully in the infeasible region, i.e. no rotation exists that can reduce the l_1 norm of the rotated quadratic coefficient tensor below $\|\mathbf{Q}\|_1 / \xi_t$. This is manifested in figure 3(a) by a sudden turn of the solution traces from the feasible region boundary upwards into the red region. This occurs because, for a given dimension M , there are only a limited number of coefficients in the tensor $\tilde{\mathbf{Q}}$ that can be shrunk to zero by any rotation of the basis functions. The threshold value increases with the ratio M/N , i.e. the l_1 norm of the rotated coefficient tensor $\tilde{\mathbf{Q}}$ can be decreased further when larger dimensions M are used. This is arguably a consequence of the fact that higher ratios M/N correspond to more degrees of freedom available to the optimiser to ensure that the sparsity constraint is satisfied. It is worth noting that in LASSO-based sparsification methods (Rubini *et al.* 2020b), the optimiser operates directly on the model coefficients and feasible solutions can always be found, with all coefficients shrunk to zero in the limiting case. However, highly sparsified models were observed to have little physical significance and poor temporal behaviour. Here, all feasible systems with varying sparsity/energy reconstruction properties are temporally stable and provide physically consistent predictions. This result is an effort to try to develop

a methodology to ensure global stability in data-driven models, extending the work done by Kaptanoglu *et al.* (2021).

In figure 3(c,b,d) the trade-off between sparsity and the energy reconstruction properties of the rotated basis is presented on the ρ - e_N plane for models constructed with $N = 20, 30, 40$, respectively. The horizontal line in each panel corresponds to the fraction of reconstructed energy of the original dense POD-Galerkin model. In these panels, squares are used to denote data for $M/N = 2$, while circles denote data for $M/N = 3$. First, it can be observed that systems for $\rho = 1$ ($\xi = 1$) do not reconstruct the entire average fluctuation kinetic energy captured by the original POD basis. This is due to the stability constraint (2.17c), producing a small rotation of the optimal POD basis so that dissipative mechanisms in the Galerkin models to ensure long-term stability are better resolved. Second, when the penalisation ξ is increased, sparser models are obtained, with higher ratios M/N enabling further reduction in density. This, however, comes at the cost of decreasing the energy optimality of the rotated basis. More importantly, the larger the model size, the more the model can be sparsified without affecting significantly the ability of the new basis to reconstruct the fluctuation kinetic energy. This suggests that the sparsification technique becomes more effective as the model complexity, and the range of scales resolved by the basis, increases. This appears to be a general trend, since similar behaviour was observed in Rubini *et al.* (2020b) using the LASSO-based *a posteriori* sparsification.

3.4. Analysis of the rotated modal structures

We now move to the analysis of the rotated spatial and temporal basis functions. We analyse a model with $N = 30$, $M/N = 3$ and density $\rho = 0.87$, obtained for $\xi = 3$ just before the solution falls into the infeasible region in figure 3(a). Figure 4(a) shows the magnitude of the entries of the rotation matrix \mathbf{X} , found from the solution of (2.17). Figure 4(b) compares the modal energies of the original POD temporal coefficients with those of the rotated modes. Figure 4(c) shows the cosine of the angle between each pair of original and rotated spatial modes, $\cos(\theta_i) = (\boldsymbol{\phi}_i, \tilde{\boldsymbol{\phi}}_i)$, which is clearly also the diagonal of \mathbf{X} , because of the orthogonality constraint (2.17d). The transformation matrix \mathbf{X} has large diagonal entries, but significant off-diagonal terms can be observed for $i, j \gtrsim 5$. This indicates that the rotated basis functions bear a strong resemblance to the original POD modes, but that the optimisation has introduced into the new basis small-scale, low-energy features to both stabilise and sparsify the rotated Galerkin model. It is observed that the high-energy temporal and spatial modes are not affected significantly by the rotation and do not differ significantly from the original POD modes. For instance, the first pair of modal energies, corresponding to the dominant fluid oscillation in the cavity, is virtually unchanged. Conversely, high-index, low-energy modes are more significantly rotated away from the corresponding original POD mode, and more significant relative differences of the modal energies are observed. We argue that this behaviour derives from the formulation of problem (2.17), constructed with the aim of generating a basis that minimises the energy loss with respect to the energetically optimal POD. As a result, the optimisation leaves mostly unchanged the most energetic modes that contribute more pronouncedly to the overall energy, and rotates by a larger extent the less energetic modes to gain in sparsity and to achieve stability.

Figures 5(a,d) show the out-of-plane vorticity component ω for the original POD spatial mode, and figures 5(b,e) the rotated mode, for indices $i = 1$ and 19, respectively. Figures 5(c,f) show the absolute value of the difference between the original and

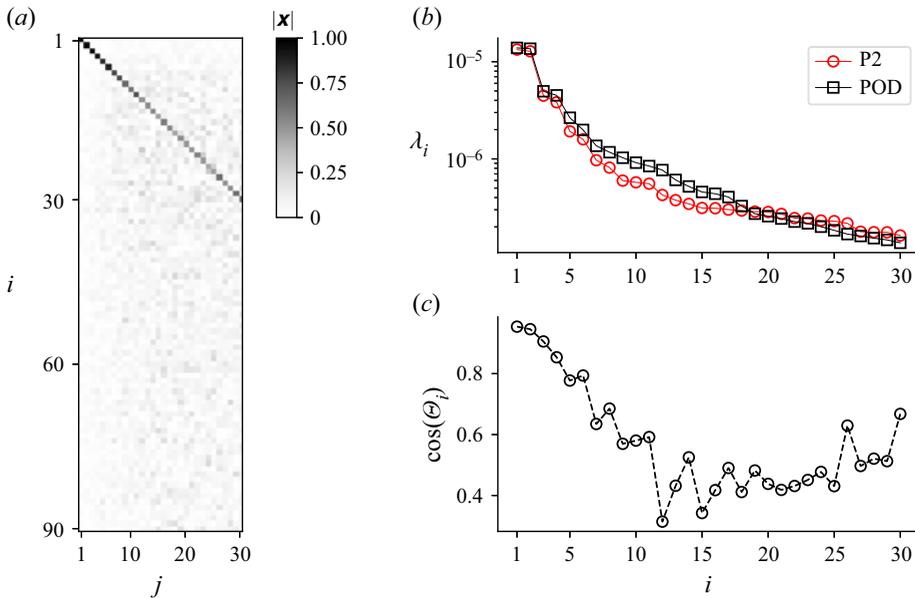


Figure 4. Optimisation results for the case $M/N = 3$, $N = 30$, $\xi = 3$ (with density $\rho = 0.87$). (a) Magnitude of the entries of the transformation matrix \mathbf{X} . (b) Distribution of the average modal energies of the original (squares) and rotated (circles) basis functions. (c) Cosine of the angle between the modes of the original and rotated basis.

rotated modes. As expected, the spatial structure of the first mode is not changed considerably by the rotation. On the other hand, mode $i = 19$ is more pronouncedly affected by the rotation, with small-scale vorticity features appearing all along the shear layer. We argue that the introduction of small-scale features is a combined effect of the stability constraint (2.17c), which enhances dissipation in the system (Balajewicz *et al.* 2016), and of the sparsity-promoting constraint (2.17b) since spatially fluctuating modes, with stronger gradients, are likely more effective to produce globally zero spatial averages involved in the projection coefficients (2.4c).

The temporal modes are affected in a similar way. This is illustrated in figures 6(a,b), showing the temporal evolution of modes $a_1(t)$ and $a_{19}(t)$ over twenty time units, and in figures 6(c,d), showing their amplitude spectra, respectively. Since the first column of \mathbf{X} is close to zero except for X_{11} , mode $a_1(t)$ and its spectral content are not affected appreciably by the rotation, except for a small general decrease of the amplitude due to decrease in energy content (see figure 4b). Conversely, the spectral content of mode $a_{19}(t)$ is remodulated by the rotation by introducing higher energy at high-frequency components, consistent with the introduction of small-scale features into the corresponding spatial mode.

3.5. Interactions identified in the sparse model

The structure of the spatial and temporal modes is only weakly modified by the optimisation, but this is sufficient to introduce sparsity in the rotated quadratic coefficient tensor $\hat{\mathbf{Q}}$ and in the rotated average triadic interaction tensor $\hat{\mathbf{N}}$ when the domain integrals (2.4c) and the temporal averages (2.8) are computed.

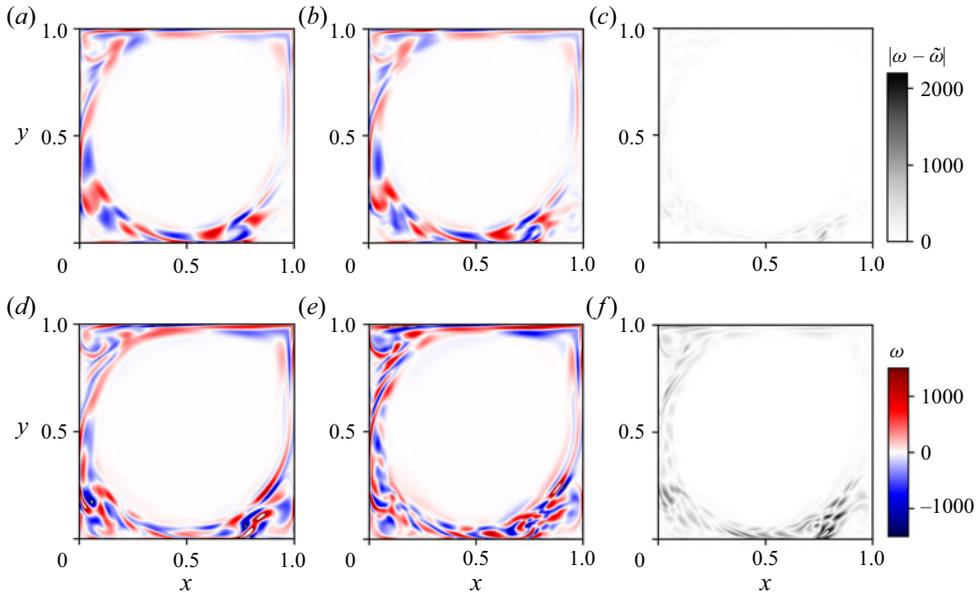


Figure 5. Vorticity fields of (a) the first original POD mode, (b) the first rotated mode (denoted as $\tilde{\omega}$), and (c) the absolute value of their difference. Panels (d–f) show the same quantities for mode $i = 19$.

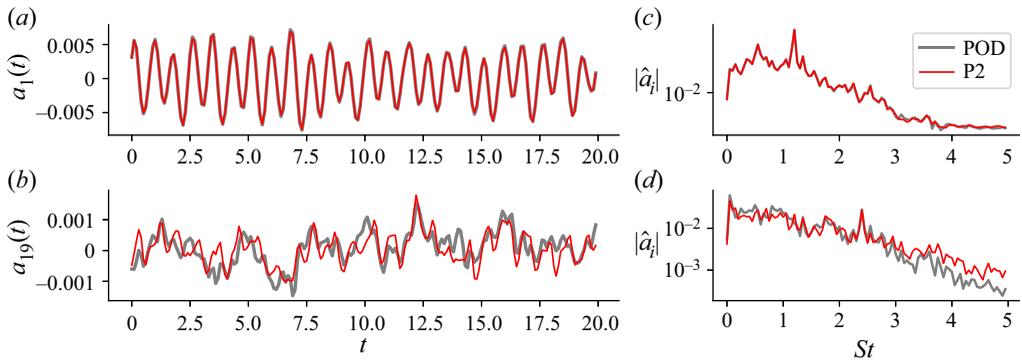


Figure 6. (a) Temporal evolution and (c) amplitude spectrum of $a_1(t)$. (b) Temporal evolution and (d) amplitude spectrum of $a_{19}(t)$. Data are reported for the original POD temporal modes and for the rotated modes of the sparse system obtained from formulation P2.

To visualise how sparsity in these tensors varies when the penalisation weight ξ is increased, we introduce the tensor ξ with entries ξ_{ijk} defined as the value of ξ in (2.17) at which the corresponding coefficient $\tilde{\mathbf{Q}}_{ijk}$ is first shrunk to zero.

Figure 7 shows two slices of ξ for $i = 1$ and $i = 30$, for a model with $N = 30$ and $M/N = 3$. Results in figures 7(a–c) are obtained with the complete formulation P2, including both the stability and sparsity-promoting constraint, while those in figures 7(d–f) are obtained with formulation P1, which includes only the sparsity-promoting constraint. Figures 7(c,f) show the modal density ρ_i as a function of the global density ρ for four modes across the hierarchy. By using formulation P2, which includes the stability constraint (2.17c), we observe that small-scale/small-scale interactions disappear first, for moderate penalisations, especially for the high-index modes, but generally across the entire

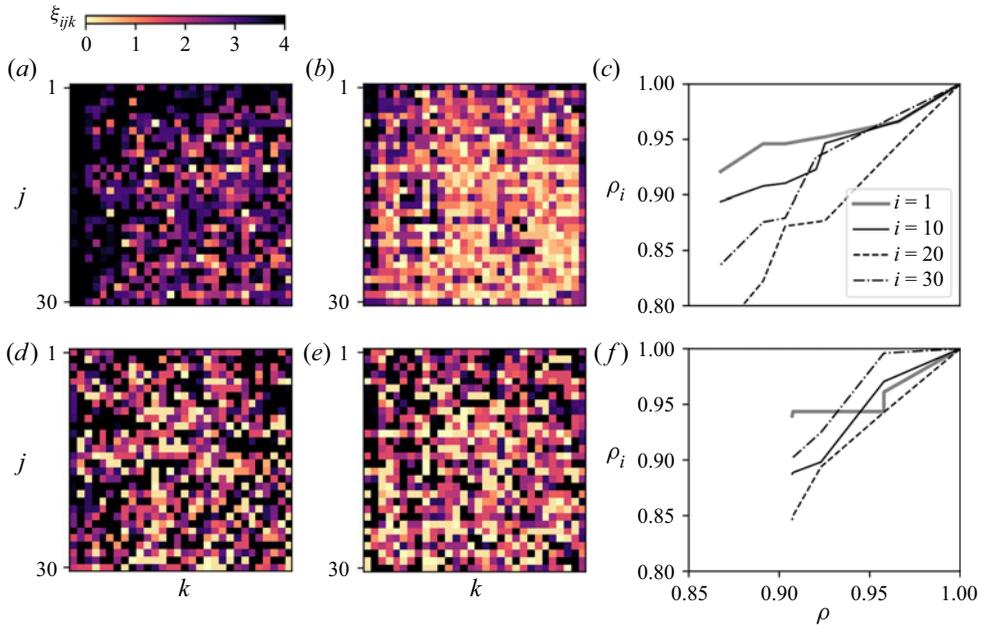


Figure 7. (a,b) Entries of the tensor ξ for $i = 1, 30$, respectively, obtained by solving problem P2. (c) Modal densities ρ_i as functions of ρ . In (d–f), the same quantities are shown from solution of the optimisation problem P1.

hierarchy of modes. Increasing the penalisation, coefficients corresponding to interactions that are local in modal space are pruned progressively, leaving only coefficients capturing non-local interactions with the low-index modes for large penalisations. The key remark is that the structure of the sparsified quadratic coefficient tensor $\tilde{\mathbf{Q}}$, solution of the proposed optimisation approach, follows the same pattern displayed by triadic energy interactions shown in figure 2(b), i.e. coefficients corresponding to energetically weak interactions are pruned first, and only relevant interactions are preserved. Interestingly, figure 7(c) indicates that high-index modes can be sparsified more efficiently. This might be related to the fact that high-index modes are also rotated more aggressively during the optimisation, to ensure that the sparsity-promoting constraint (2.17b) is satisfied and with minor effect of the overall energy reconstruction ability.

Nonetheless, considering now the solution obtained from formulation P1 in figures 7(d–f), it is clear that many quadratic coefficients are indeed shrunk to zero during the optimisation, and similar global densities are obtained. However, the sparsity pattern does not have a clear relation with the original structure of energy interactions. By contrast, coefficients corresponding to important energy interactions have been shrunk to zero since the optimisation problem P1 is driven entirely by the sparsity-promoting constraint. While in both formulations the same penalisation on the l_1 norm of the rotated tensor $\tilde{\mathbf{Q}}$ is used, formulation P1 lacks any information regarding the dynamics and temporal evolution of the Galerkin model and the structure of intermodal energy transfers. This results in an unphysical equal contraction of all quadratic coefficients, across all modes and independently of the strength of the energy interactions that they represent. The effect of this behaviour can also be noticed in the ρ_i – ρ curves in figure 7(f), where different individual modes are sparsified by a similar amount. Nevertheless, no significant differences are observed between models obtained from the two formulations on the

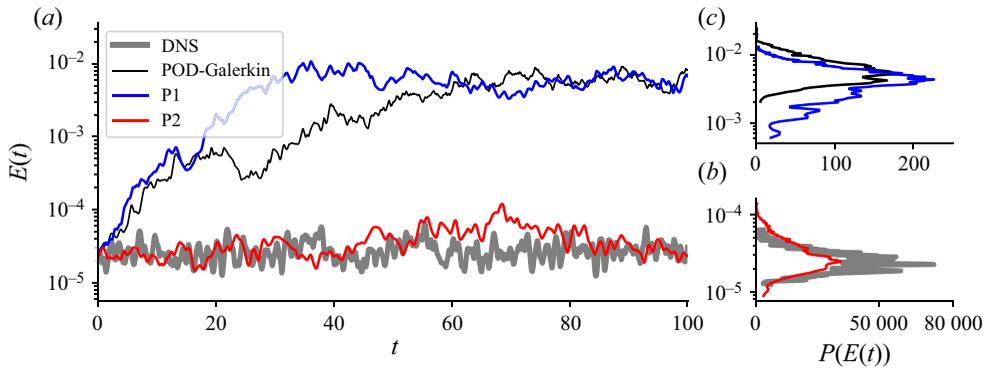


Figure 8. (a) Temporal evolution of the integral fluctuation kinetic energy $E(t)$. The temporal performances of the two sparse models obtained by solving problem (2.14) are compared against DNS and the dense model obtained with Galerkin projection. Panels (b,c) show the probability distribution of the energy $E(t)$, obtained for long-time integration of the Galerkin models.

ρ - e_N plane. From a physical standpoint, this suggests that the proposed approach can produce sparse models with realistic physics only if some notion of the flow dynamics is included in the optimisation procedure. In other words, only targeting the sparsity of the rotated coefficient tensor is not conducive to realistic flow models. In the present case, this is achieved through constraint (2.17c). This, however, ensures only that the model reproduces that same average fluctuation energy observed in DNS. Interestingly, this global metric is sufficient to capture correctly the local nature of intermodal energy transfers.

3.6. Temporal integration and energy analysis of the sparsified system

In this section, we analyse the temporal behaviour of the sparse reduced-order models obtained by solving problems P1 and P2, and consider their rotated triadic interactions tensors $\tilde{\mathbf{N}}$ and the average energy budget of (2.7). The same minimum-density configuration studied in previous sections, with $M/N = 3$, $N = 30$ and $\xi = 3$, is considered. This configuration is used here as a representative example of the sparse models obtained with the proposed approach. In fact, the spatio-temporal statistical behaviour of these models is affected only weakly by the regularisation parameter ξ , but is mostly dominated by the inclusion of the stability constraint, regardless of the sparsity. Models are integrated forward in time with an implicit time-stepping scheme for $T = 500$ time units, from an initial condition obtained from one of the snapshots. Figure 8(a) shows the first one hundred time units of the temporal evolution of the turbulent kinetic energy (2.5) for these two models, compared with the evolution from DNS and from the original POD-Galerkin model. Figures 8(b,c) show the probability density function of the same quantity, computed over a longer time span. As expected, the deficit of dissipation in the original POD-Galerkin model produces fluctuation kinetic energy levels approximately two orders of magnitude larger than the reference value from DNS. This behaviour is well known (see e.g. Östth *et al.* 2014; Noack *et al.* 2016).

The sparse Galerkin model obtained without the stability constraint (problem P1) also reproduces this behaviour. Interestingly, the sparse model obtained from this formulation spans a much wider range of flow scales than the POD model ($N = 30$ modes are rotated within a space of $M = 90$ dimensions), but it still displays dynamical behaviour that is

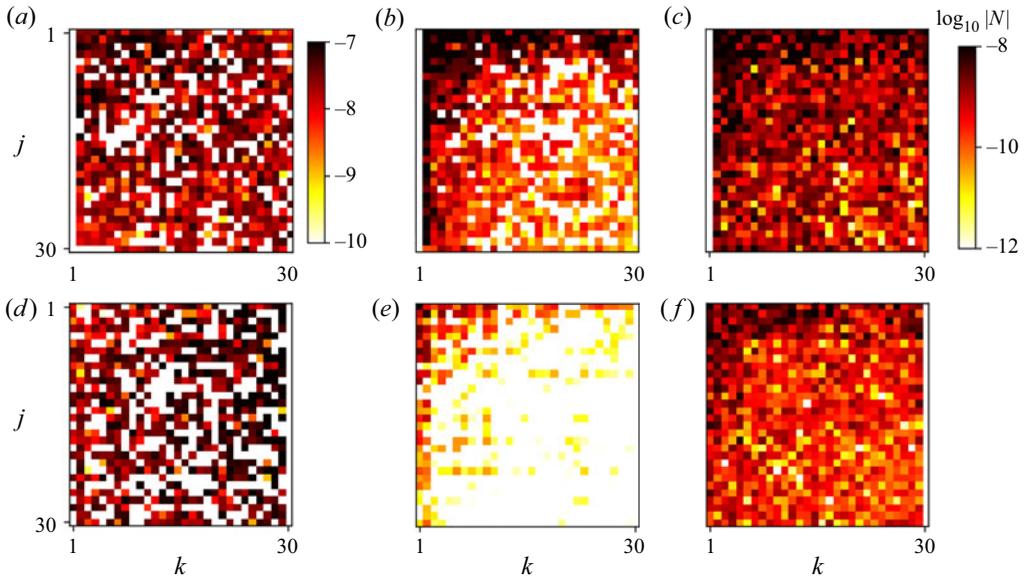


Figure 9. (a–c) Intensity of the average nonlinear energy transfer rate \tilde{N} for mode $i = 1$ in models obtained from projection and formulations P2 and P1, respectively. Panels (d–f) show the same quantity for mode $i = 30$.

usually attributed to a lack of viscous dissipation. This shows that optimising for sparsity alone is not sufficient to obtain physically realistic models, because energy transfers from energy-producing eddies and dissipative structures are not necessarily captured correctly during the optimisation. Conversely, the sparsified and stabilised model obtained from problem P2 has realistic temporal behaviour and resolves correctly the average energy and its fluctuation observed in DNS. This suggests that including in the optimisation information on how modal structures are supposed to interact with each other to satisfy the overall power budget is key to obtaining sparse Galerkin models with adequate predictive ability. Note that the long-term performance of the models obtained in the present work is generally superior to that of models sparsified *a posteriori*, using a LASSO-based approach (Rubini *et al.* 2020b; Rubini, Lasagna & Da Ronch 2020a). More specifically, LASSO-based models have been found to be temporally accurate over a time span comparable to that of the data used for the sparsification. Conversely, the present *a priori* sparsified systems are temporally stable for much longer time horizons. In addition, numerical solutions converge to the asymptotic attractor for much larger perturbations of the initial conditions. We argue that this is likely the consequence of enforcing a stronger consistency between the modal structures and the corresponding Galerkin model.

The structure of the nonlinear energy transfer rate tensor \mathbf{N} for the rotated Galerkin models is reported in figures 9(a–c) and 9(d–f), showing two slices for $i = 1, 30$, respectively. Figures 9(a,d) refer to the model obtained from Galerkin projection onto the original POD subspace and will be used as reference. Here, the temporal coefficients are the projections onto the DNS snapshots. Data for the model obtained from problem P2 are shown in figures 9(b,e), while data for the model obtained with the P1 formulation are shown in figures 9(c,f). For these models, temporal coefficients are obtained from forward integration of the Galerkin models. It can be observed that the model from formulation P2 displays a pattern of interactions consistent with the DNS data projected onto the original POD modes, in terms of both organisation across modes as well as strength.

The organisation of the interactions is similar to that in the dense Galerkin model obtained using the *a priori* stabilisation method of Balajewicz *et al.* (2013), although a more aggressive reduction in the strength of the interactions between the high-index modes is observed. This might be a physical effect (small scales tend to be advected primarily by the large ones), but a second possible explanation is that this is the product of the optimisation procedure adopted in this work. To meet the sparsity constraint with minimum impact on the energy reconstruction ability of the rotated basis, the optimiser rotates the high-index modes more aggressively, causing higher sparsity in the dynamical equations of modes associated with small scales. By contrast, energy interactions in the model obtained from the formulation P1 are orders of magnitude more intense, across all triads. This is a manifestation of the lack of dissipation and the consequential over-prediction of energy, across all modes. Interestingly, it can be observed that the HL/LH asymmetry observed in figure 2 is preserved throughout the sparsification. Conversely, this physical feature is lost in the *a posteriori* approach (Rubini *et al.* 2020b) for numerical reasons. It is argued that the *a priori* approach is preferable for systems where the tensor \mathbf{N} has a non-trivial structure, because the link between modal structures and energy interaction tensor is fully preserved.

To conclude this section, we analyse the energy budget of the original and rotated Galerkin models, examining the linear and quadratic energy transfer terms in the average energy budget equation (2.7) on a mode-by-mode basis (Noack *et al.* 2008). To characterise the total nonlinear energy transfer, we use the quantity

$$T_i = \sum_{j=1}^N \sum_{k=1}^N Q_{ijk} \overline{a_i a_j a_k}, \quad i = 1, \dots, N, \quad (3.2)$$

satisfying $\sum_i T_i = 0$ when the quadratic term conserves energy. In addition, the linear energy transfer term is decomposed into its three constitutive components and we use $L^{(m)}$ for the m th term of the definition (2.4b). These are, in order, the viscous dissipation rate and the production and convection rates mediated by ‘macroscopic’ interactions with the mean flow. The average energy flows associated with these terms are displayed in figure 10 as a function of the modal index. Data for three models are reported: the original POD-Galerkin model (figures 10a,d,g,j,m) with temporal modes obtained from projection on the DNS data, the model obtained from the complete formulation P2 at the minimum density (figures 10b,e,h,k,n), and the model obtained from formulation P1 (figures 10c,f,i,l,o). For these two models, temporal coefficients are obtained from long-time integration. All models are constructed with $N = 30$ and, where applicable, $M/N = 3$.

Figures 10(a–i) show the viscous, production and convection components of the linear energy flow term, while their sum is reported in figures 10(j–l). The nonlinear transfer term is reported in figures 10(m–o). It can be observed that all three models describe viscous dissipation effects in a quantitatively similar manner, and that convection is the only other relevant term of the linear energy flow, while the production term is negligible. This result can be explained by noting that the advection of small-scale features on the shear layer bounding the primary vortical structure operated by the mean recirculatory motion is a significant dynamical feature of the lid-driven cavity flow. However, the model obtained with the P1 formulation, where the rotation of the original POD basis is driven only by the requirement to satisfy the sparsity constraint, results in higher average production rates across the entire spectrum of modes. Overall, the linear energy flow term for the original POD-Galerkin model (figure 10j) exhibits a positive net energy flow for the first

A priori sparsification of Galerkin models

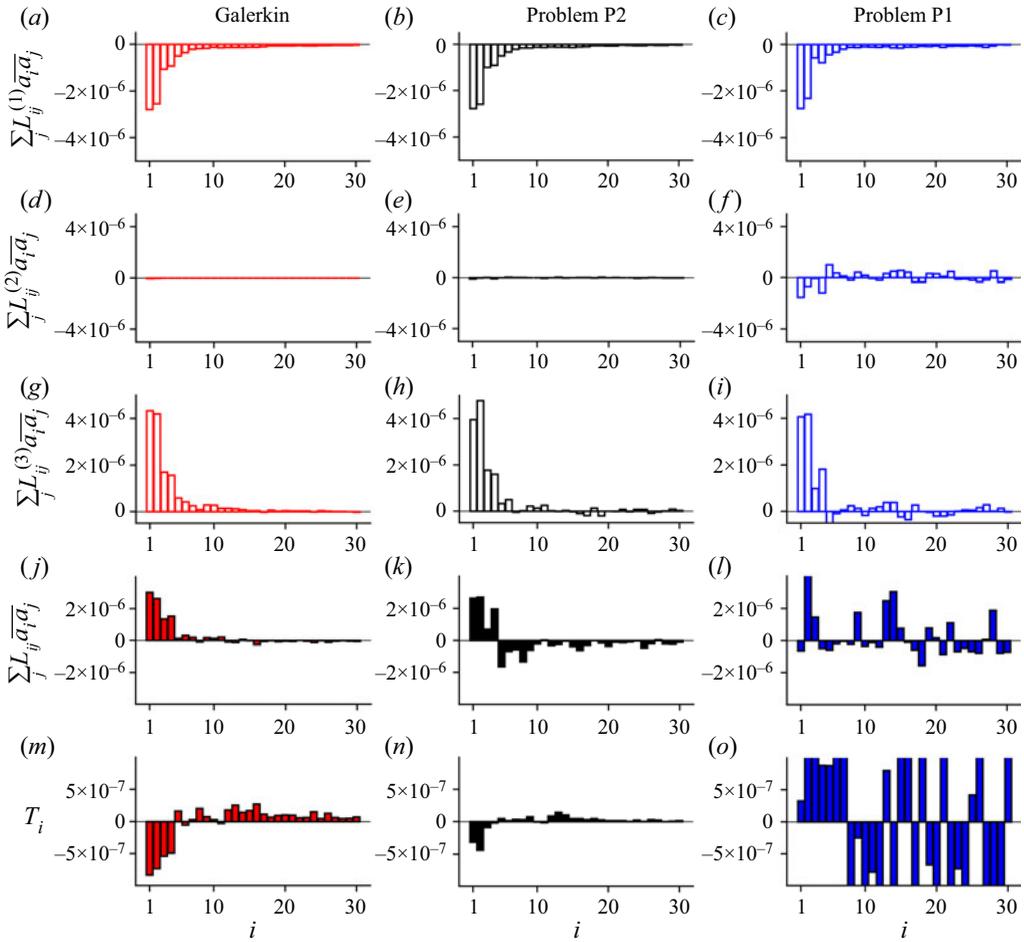


Figure 10. Distributions of time-averaged energy flows as functions of the modal index for models with $N = 30$. Data are reported for the original POD-Galerkin model, (a,d,g,j,m), and for the rotated models from formulations P2 and P1, (b,e,h,k,n) and (c,f,i,l,o), respectively. Panels (a–i) include the energy flows associated to the three components of the linear coefficient tensor (2.4b), corresponding to the viscous dissipation rate and the production and convection energy flow mediated by the mean flow. The sum of the linear energy flows is reported in (j–l). The average nonlinear transfer rates are reported in (m–o).

four modes, corresponding to the dominant structure and flow oscillation in the cavity, but only moderate dissipation in the remaining modes. This results in an imbalance between extraction of energy from the mean flow and dissipation, leading to the over-prediction of kinetic energy as shown in figure 8(a). The mean nonlinear transfer rate T_i shown in figure 10(m) displays the correct transfer direction, with energy injected from the large scales (negative transfer rate for low-index modes) to the small scales (positive rate for the remaining modes). The model obtained from the solution of the complete formulation P2 (figure 10k) displays a better balance, since a negative net energy flow rate is observed for $i > 4$, due to effect of the stability constraint (2.17b). The distribution and direction of the nonlinear transfers is also well preserved, although lower transfer rates are observed for all modes. This results from the combined effect of the sparsity-promoting constraint, which reduces the magnitude of the quadratic coefficients, and the reduction of the modal energies, due to the loss of optimality from the POD. On the other hand, energy transfers in the model obtained from formulation P1, which does not include any

dynamical information on the fluctuation kinetic energy budget (2.7), have a significantly different structure. In fact, both the linear and nonlinear transfer rates show a marked increase, and the organisation of energy flows across the hierarchy of modes has lost the original physical character.

These results show that to obtain a sparse model that is consistent with the flow physics observed in direct simulation, it is of paramount importance to retain in the sparsification algorithm some information regarding the temporal dynamics. This is consistent with what was observed by Loiseau & Brunton (2018), who showed that an additional constraint on the energy conservation is one way to enhance the temporal accuracy of a fluid model reconstructed from data.

4. Conclusions

Scale interactions in turbulent flows are sparse, and the motion at any given length scale depends most prominently on the dynamics of a subset of all other scales. In addition to the recent work of Schmidt (2020), this paper is one attempt at developing a model order-reduction technique that leverages this fact, so that the analysis and interpretation of scale interactions in complex flows is facilitated. Here, we have proposed an *a priori* sparsification approach, whereby a set of basis functions describing coherent structures that interact minimally with one another is sought for. Scale interaction sparsity is then defined by the sparsity of the quadratic coefficient tensor of the Galerkin model constructed from projection on this basis. As opposed to our previous sparsification approach (Rubini *et al.* 2020*b*), where model coefficients corresponding to weak nonlinear interactions between a predetermined set of basis functions are pruned *a posteriori* using l_1 regression, the present methodology maintains the exact link between the (sparse) Galerkin model and the modal structures utilised for the projection.

In practice, a set of N basis functions is expressed as a rotation of N proper orthogonal decomposition (POD) modes within a larger POD subspace (of dimension $M > N$), with the idea of altering energy paths across the model by rotating the subspace utilised for projection whilst minimising the loss of energy representation ability of the original POD basis. The rotation matrix is then found from the solution of a constrained optimisation problem where (i) the energy loss with respect to the original POD basis is minimised, and (ii) an inequality constraint involving the l_1 norm of the quadratic coefficient tensor associated to the rotated basis promotes sparsity. This formulation is augmented with a further constraint that enforces long-term temporal stability, following the approach of Balajewicz *et al.* (2016). This optimisation problem depends on a single user-controllable parameter controlling the l_1 norm of the rotated quadratic coefficient tensor and the trade-off between energy reconstruction and sparsity.

To demonstrate the approach, we considered the incompressible two-dimensional lid-driven cavity flow at Reynolds number $Re = 2 \times 10^4$, where the motion is chaotic and energy interactions are scattered in modal space. Results show that there is competition between sparsification and energy representation, and that this trade-off depends on the model size N and the ratio M/N . In particular, larger models can be sparsified more aggressively with less impact on the energy reconstruction ability. On the other hand, higher sparsity can be obtained by rotating the same N POD basis functions within larger subspaces (a higher ratio M/N), since the increased flexibility allows shrinking to zero a larger fraction of quadratic interaction coefficients. A deeper analysis of energy paths shows that the distribution of intermodal energy transfers in the rotated model is qualitatively similar to that of the original POD-Galerkin system. More specifically, coefficients of the quadratic interaction tensor corresponding to

large-scale/large-scale and large-scale/small-scale interactions are preserved, but those defining weaker small-scale/small-scale interactions are shrunk to zero during the optimisation. This result is in agreement with the established picture of triadic interactions in two-dimensional flows and with previous results of sparsification of reduced-order models (Rubini *et al.* 2020*b*). More interestingly, we observed that a physically consistent organisation of the interactions and a stable long-term behaviour can be obtained only by augmenting the sparsification procedure with the temporal stability constraint. In fact, models obtained without such a constraint inherit well-known temporal stability issues displayed by dense POD-Galerkin models.

Some aspects deserve further discussion. Unlike in l_1 -regression-based sparsification methods (Loiseau & Brunton 2018; Rubini *et al.* 2020*b*), where the model coefficients are the optimisation variables, here these coefficients are cubic polynomial functions of the optimisation variables, the entries of the rotation matrix. The first consequence is that the optimisation problem is non-convex, and many local optima, i.e. many different sets of basis functions, exist. It was observed that initial guesses close to the original POD basis, i.e. with good energy representation ability, converge repeatedly to the same optimal solution, which has a consistent physical interpretation. Arguably, initial guesses lying farther from the POD might lead to sparser solutions at the cost of a worse energy reconstruction capability. However, the lack of strong uniqueness guarantees (as for many other modal decomposition techniques) may render the approach questionable. Second, it is not possible to shrink to zero an arbitrary number of model coefficients, as the strong link between the basis functions and the Galerkin model must always be maintained. As a result, the present *a priori* sparsification technique produces relatively denser Galerkin models than the *a posteriori* LASSO-based approach considered in Rubini *et al.* (2020*b*) for the same test case. Rotating the original POD basis into higher-dimensional subspaces, with ratio M/N higher than that considered in the present work, to further increase sparsity is possible, although it would necessarily result in increased computational costs for the optimisation.

Funding. The authors gratefully acknowledge support for this work from the Air Force Office of Scientific Research (Grant no. FA9550-17-1-0324, Program Manager Dr D. Smith).

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

Author ORCIDs.

 Riccardo Rubini <https://orcid.org/0000-0003-1883-5236>;

 Davide Lasagna <https://orcid.org/0000-0002-6501-6041>;

 Andrea Da Ronch <https://orcid.org/0000-0001-7428-6935>.

Appendix A. Sparsity-promoting effect of the l_1 -based constraint

In this appendix we visualise how the l_1 norm appearing in the constraint (2.17*b*) favours sparse solutions. For simplicity, we use a small ($N = 3$, $M = 5$) reduced-order model of the lid-driven cavity flow considered in the previous sections. A more formal discussion based on the proximity operator theory can also be formulated, using the classical LASSO formulation as a starting point (Friedman *et al.* 2008). In figure 11(*a*), the ratio $\|\tilde{\mathbf{Q}}\|_1/\|\mathbf{Q}\|_1$ is shown as a function of the entries X_{23} and X_{12} of the rotation matrix. When these two parameters are varied, all the other off-diagonal entries of \mathbf{X} are set to zero. Several sharp valleys can be observed, arising from the non-smooth nature of the l_1 norm. Note that $\|\tilde{\mathbf{Q}}\|_1/\|\mathbf{Q}\|_1$ is a non-convex cubic function of the entries of the rotation matrix, and several local minima can be identified. Figure 11(*b*) shows a cut of figure 11(*a*)

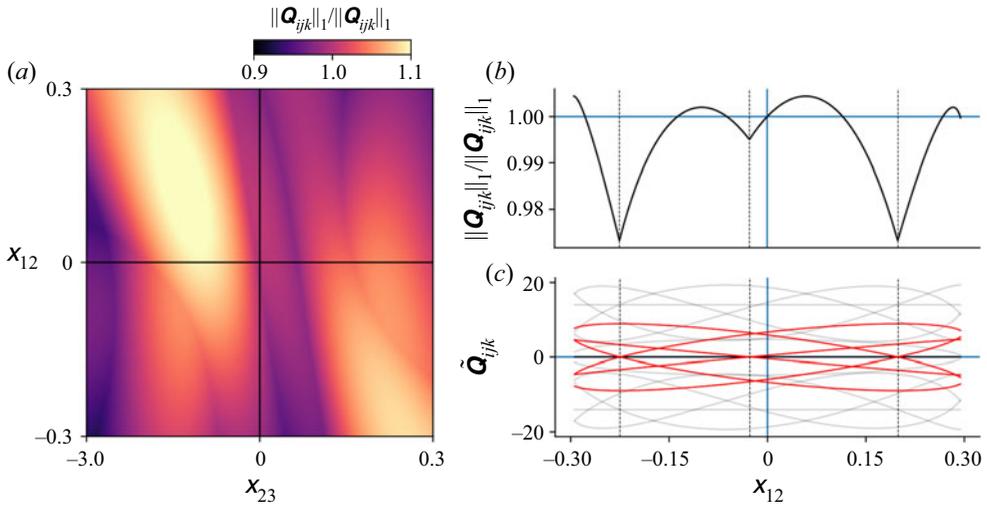


Figure 11. (a) Map of $\|\tilde{\mathbf{Q}}\|_1/\|\mathbf{Q}\|_1$ projected on the X_{12} – X_{23} plane. (b) One-dimensional cut of the same quantity along the coordinate X_{12} . (c) The twenty-seven entries of the rotated $\tilde{\mathbf{Q}}$ as a function of X_{12} . The six entries that vanish identically at one of the sharp points of $\|\tilde{\mathbf{Q}}\|_1$ are highlighted in red.

along the X_{12} coordinate, where minima are identified by the dashed vertical lines. Each minimum corresponds to a point where one of the entries of $\tilde{\mathbf{Q}}_{ijk}$ crosses the zero axis, as illustrated in figure 11(c). During the optimisation, the sparsity-promoting constraint pushes the solution towards one of these valleys, due to the strong gradient at the point of non-differentiability, resulting in a sparse coefficient tensor $\tilde{\mathbf{Q}}$. However, unlike in the *a posteriori* LASSO-based sparsification methods (Brunton *et al.* 2016; Rubini *et al.* 2020b), not all quadratic coefficients can be set to zero simultaneously by an arbitrary rotation, as is clear from figure 11(c). In fact, the tensor $\tilde{\mathbf{Q}}$ depends nonlinearly on the rotation \mathbf{X} , while in LASSO-based methods the tensor coefficients are directly the optimisation variables of the problem.

Appendix B. Dependence of the solution from the initial condition

Problem (2.17) is non-convex due to the constraint on the l_1 norm of the rotated quadratic interaction coefficient tensor $\tilde{\mathbf{Q}}$.

To understand how this feature affects the results of § 3, we consider in this appendix different sets of initial guesses constructed as increasingly larger perturbations of the identity as

$$\mathbf{X} = \begin{bmatrix} I_{N \times N} \\ \mathbf{0} \end{bmatrix} + \epsilon \mathbf{R}, \tag{B1}$$

with $\mathbf{R} \in \text{Re}^{M \times N}$ a randomly generated rotation matrix, satisfying $\mathbf{R}^T \mathbf{R} = I_{N \times N}$. We generated ten random rotation matrices for five values of ϵ in the range $[10^{-8}, 10]$ and solved problem (2.17) by keeping $\xi = 2$ to target an intermediate density in the feasibility region, using $N = 30$ and $M = 90$.

Results are shown in figure 12. Figure 12(a) quantifies the effects of ϵ on the average rotation angle θ between columns of the initial guess and those of the matrix $[I_{N \times N}, \mathbf{0}]^T$. Circles and vertical bars identify the average and standard deviation of this quantity across the ten different samples. In figures 12(b,c), the largest modal energy $\tilde{\lambda}_1$ of the rotated

A priori sparsification of Galerkin models

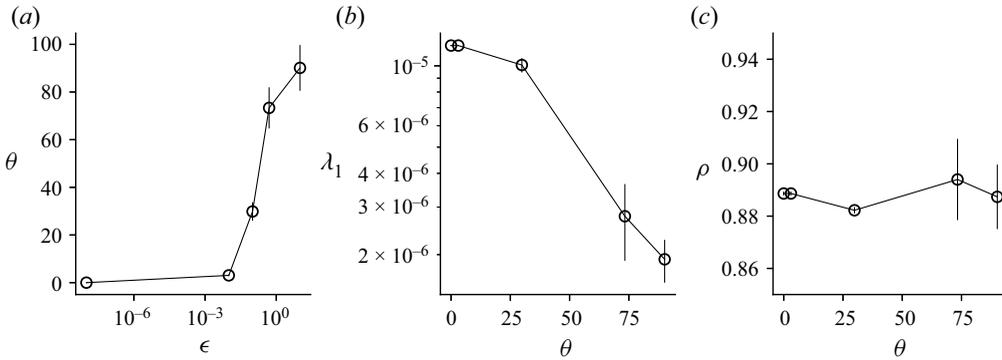


Figure 12. Effect of the initial guess on the optimisation results. (a) Relation between the initial guess perturbation size ϵ and the average rotation angle θ away from the original POD basis. Mean and standard deviation of (b) the first modal energy $\tilde{\lambda}_1$, and (c) density ρ , from rotated systems obtained from initial guesses of increasing distance from the original POD basis.

temporal coefficients and the density ρ of the optimal Galerkin model, respectively, are shown as functions of the average θ . Results show that for small perturbations of the original POD basis, all initial guesses converge to the same solution, since $\tilde{\lambda}_1$ and ρ are the same (up to the tolerances set in the optimisation). Conversely, initial guesses corresponding to larger (random) rotation of the original POD basis converge to different solutions, corresponding to modal structures with lower reconstruction ability, but similar model density.

Appendix C. Optimisation of tensor operations and scaling of costs

The sparsity-promoting constraint involves polynomial functions of the optimisation variables. An analytical expression of the gradient of this constraint with respect to the rotation \mathbf{X} can also be obtained, enabling fast gradient-based optimisation to be utilised. In addition, a careful examination of the tensor operations involved in the computation of the constraint and its gradient shows that a significant reduction of the scaling of costs can be obtained by reorganising some tensorial computations and using intermediate temporary variables (see Pfeifer, Haegeman & Verstraete 2014). For instance, a naive calculation of all entries of the rotated tensor (2.11), required in the evaluation of the sparsity-promoting constraint (2.14b), takes $O(M^3N^3)$ operations since six different nested for loops (one for each index) are involved in total. However, a careful examination of (2.11) shows that the rotated tensor can also be obtained by first introducing the auxiliary temporary tensors $\mathbf{A} \in \text{Re}^{M \times M \times N}$ and $\mathbf{B} \in \text{Re}^{M \times N \times N}$, and then computing

$$A_{qri} = \sum_{p=1}^M Q_{pqr} X_{pi}, \quad (\text{C1a})$$

$$B_{rij} = \sum_{r=1}^M A_{qri} X_{qj}, \quad (\text{C1b})$$

$$\tilde{Q}_{ijk} = \sum_{q=1}^M B_{rij} X_{rk}, \quad (\text{C1c})$$

with $i, j, k = 1, \dots, N$ and $q, r = 1, \dots, M$. Collectively, these steps take only $O(M^3N) + O(M^2N^2) + O(MN^3)$ operations, as can be inferred from the summation indices.

On the other hand, the gradient of the sparsity-promoting constraint with respect to the entries X_{mn} of the rotation matrix \mathbf{X} requires the computation of

$$\frac{\partial \|\tilde{\mathbf{Q}}\|_1}{\partial X_{mn}} = \sum_{i,j,k=1}^N \frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} \text{sign}(\tilde{Q}_{ijk}), \quad (\text{C2})$$

where the gradient $\partial \tilde{Q}_{ijk} / \partial X_{mn}$ can be expanded using Kronecker's delta δ_{ij} as

$$\frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} = \sum_{p,q,r=1}^M Q_{pqr} (X_{pi} X_{qj} \delta_{kn} \delta_{rm} + X_{pi} X_{rk} \delta_{jn} \delta_{qm} + X_{rk} X_{qj} \delta_{in} \delta_{pm}). \quad (\text{C3})$$

Overall, a naive computation of (C3) would require $O(M^4N^4)$ operations, since the tensor expressions involve a total of eight indices. A reorganisation of the operations involved in (C3) similar to that described above leads to tensor operations over four tensor indices only, with computational costs scaling similarly to the rotation.

Lastly, we underline that in the present implementation, we chose to store the whole tensors (2.11), (C3) and (C2) in memory. This approach is computationally efficient since it enables us to use optimised packages for tensorial calculus, but it is very expensive memory-wise. A further optimisation can be obtained by assembling (2.11) and (C2) without fully storing of the intermediate steps.

REFERENCES

- AMSALLEM, D. & FARHAT, C. 2012 Stabilization of projection-based reduced-order models. *Intl J. Numer. Meth. Engng* **91** (4), 358–377.
- ARBABI, H. & MEZIĆ, I. 2017 Study of dynamics in post-transient flows using Koopman mode decomposition. *Phys. Rev. Fluids* **2** (12), 124402.
- AUTERI, F., PAROLINI, N. & QUARTAPELLE, L. 2002 Numerical investigation on the stability of singular driven cavity flow. *J. Comput. Phys.* **183** (1), 1–25.
- BALAJEWICZ, M., DOWELL, E. & NOACK, B. 2013 Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier–Stokes equation. *J. Fluid Mech.* **729**, 285–308.
- BALAJEWICZ, M., TEZAUR, I. & DOWELL, E. 2016 Minimal subspace rotation on the Stiefel manifold for stabilization and enhancement of projection-based reduced order models for the compressible Navier–Stokes equations. *J. Comput. Phys.* **321**, 224–241.
- BRASSEUR, J. & WEI, C. 1994 Interscale dynamics and local isotropy in high Reynolds number turbulence within triadic interactions. *Phys. Fluids* **6**, 842–870.
- BRUNTON, S.L., PROCTOR, J.L. & KUTZ, J.N. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci.* **113** (15), 3932–3937.
- CAZEMIER, W., VERSTAPPEN, R.W.C.P. & VELDMAN, A.E.P. 1998 Proper orthogonal decomposition and low-dimensional models for driven cavity flows. *Phys. Fluids* **10** (7), 1685–1699.
- CORDIER, L., EL MAJD, B.A. & FAVIER, J. 2010 Calibration of POD reduced-order models using Tikhonov regularization. *Intl J. Numer. Meth. Fluids* **63** (2), 269–296.
- COUPLET, M., SAGAUT, P. & BASDEVANT, C. 2003 Intermodal energy transfers in a Proper Orthogonal Decomposition–Galerkin representation of a turbulent separated flow. *J. Fluid Mech.* **491**, 275–284.
- CRAIK, A. 1971 Non-linear resonant instability in boundary layers. *J. Fluid Mech.* **50** (2), 393–413.
- DESHMUKH, R., MCNAMARA, J.J., LIANG, Z., KOLTER, J.Z. & GOGULAPATI, A. 2016 Model order reduction using sparse coding exemplified for the lid-driven cavity. *J. Fluid Mech.* **808**, 189–223.
- DOMARADZKI, J.A., LIU, W., HÄRTEL, C. & KLEISER, L. 1994 Energy transfer in numerically simulated wall-bounded turbulent flows. *Phys. Fluids* **6** (4), 1583–1599.
- ERICHSON, N.B., ZHENG, P., MANOHAR, K., BRUNTON, S.L., KUTZ, J.N. & ARAVKIN, A.Y. 2020 Sparse principal component analysis via variable projection. *SIAM J. Appl. Maths* **80** (2), 977–1002.

- FICK, L., MADAY, Y., PATERA, A.T. & TADDEI, T. 2018 A stabilized POD model for turbulent flows over a range of Reynolds numbers: optimal parameter sampling and constrained projection. *J. Comput. Phys.* **371**, 214–243.
- FRIEDMAN, J., HASTIE, T. & TIBSHIRANI, R.J. 2008 *The Elements of Statistical Learning*. Springer.
- GALLETTI, B., BRUNEAU, C.H., ZANNETTI, L. & IOLLO, A. 2004 Low-order modelling of laminar flow regimes past a confined square cylinder. *J. Fluid Mech.* **503**, 161–170.
- GRIMBERG, S., FARHAT, C & YOUKILIS, N. 2020 On the stability of projection-based model order reduction for convection-dominated laminar and turbulent flows. *J. Comput. Phys.* **419**, 109681.
- JIN, B., SYMON, S. & ILLINGWORTH, S.J. 2021 Energy transfer mechanisms and resolvent analysis in the cylinder wake. *Phys. Rev. Fluids* **6** (2), 024702.
- JOHNSON, S.G. 2014 The nlopt nonlinear-optimization package. <http://github.com/stevengj/nlopt>.
- JOLLIFFE, I.T., TRENDAFILOV, N.T. & UDDIN, M. 2003 A modified principal component technique based on the LASSO. *J. Comput. Graph. Stat.* **12** (3), 531–547.
- JOVANOVIĆ, M.R., SCHMID, P.J. & NICHOLS, J.W. 2014 Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **26** (2), 024103.
- KAPTANOGLU, A.A., CALLAHAM, J.L., HANSEN, C.J., ARAVKIN, A. & BRUNTON, S.L. 2021 Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Phys. Rev. Fluids* **6** (9), 094401.
- KOLMOGOROV, A.N. 1991 The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Proc. R. Soc. Lond. A* **434** (1890), 9–13.
- KRAICHNAN, R.H. 1971 Inertial-range transfer in two- and three-dimensional turbulence. *J. Fluid Mech.* **47** (3), 525–535.
- LAVAL, J.P., DUBRULLE, B. & NAZARENKO, S. 1999 Nonlocality of interaction of scales in the dynamics of 2D incompressible fluids. *Phys. Rev. Lett.* **83** (20), 4061.
- LEE, M.W. & DOWELL, E.H. 2020 Improving the predictable accuracy of fluid Galerkin reduced-order models using two POD bases. *Nonlinear Dyn.* **101** (2), 1457–1471.
- LOISEAU, J.C. & BRUNTON, S.L. 2018 Constrained sparse Galerkin regression. *J. Fluid Mech.* **838**, 42–67.
- MOFFATT, H.K. 2014 Note on the triad interactions of homogeneous turbulence. *J. Fluid Mech.* **741**, R3.
- NOACK, B., MORZYŃSKI, M. & TADMOR, G. 2011 *Reduced-Order Modelling for Flow Control*, vol. 528. Springer Science & Business Media.
- NOACK, B.R., PAPAS, P. & MONKEWITZ, P.A. 2005 The need for a pressure-term representation in empirical Galerkin models of incompressible shear flows. *J. Fluid Mech.* **523**, 339–365.
- NOACK, B.R., SCHLEGEL, M., AHLBORN, B., MUTSCHKE, G., MORZYŃSKI, M., COMTE, P. & TADMOR, G. 2008 A finite-time thermodynamics of unsteady fluid flows. *J. Non-Equilib. Thermodyn.* **33** (2), 103–148.
- NOACK, B.R., STANKIEWICZ, W., MORZYŃSKI, M. & SCHMID, P.J. 2016 Recursive dynamic mode decomposition of transient and post-transient wake flows. *J. Fluid Mech.* **809**, 843–872.
- OHKITANI, K. 1990 Nonlocality in a forced two-dimensional turbulence. *Phys. Fluids A* **2** (9), 1529–1531.
- OLSHAUSEN, B.A. & FIELD, D.J. 1996 Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381** (6583), 607–609.
- ÖSTH, J., NOACK, B.R., KRAJNOVIĆ, S., BARROS, D. & BORÉE, J. 2014 On the need for a nonlinear subscale turbulence term in POD models as exemplified for a high-Reynolds-number flow over an Ahmed body. *J. Fluid Mech.* **747**, 518–544.
- PENG, Y.H., SHIAU, Y.H. & HWANG, R.R. 2003 Transition in a 2-D lid-driven cavity flow. *Comput. Fluids* **32** (3), 337–352.
- PFEIFER, R.N.C., HAEGEMAN, J. & VERSTRAETE, F. 2014 Faster identification of optimal contraction sequences for tensor networks. *Phys. Rev. E* **90**, 033315.
- POPE, S.B. 2001 *Turbulent Flows*. Cambridge University Press.
- REMPFER, D. & FASEL, H.F. 1994a Dynamics of three-dimensional coherent structures in a flat-plate boundary layer. *J. Fluid Mech.* **275**, 257–283.
- REMPFER, D. & FASEL, H.F. 1994b Evolution of three-dimensional coherent structures in a flat-plate boundary layer. *J. Fluid Mech.* **260**, 351–375.
- ROWLEY, C.W. & DAWSON, S.T.M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.
- RUBINI, R., LASAGNA, D. & DA RONCH, A. 2020a I1-based calibration of POD-Galerkin models of two-dimensional unsteady flows. *Chin. J. Aeronaut.* **34** (1), 226–236.
- RUBINI, R., LASAGNA, D. & DA RONCH, A. 2020b The I1-based sparsification of energy interactions in unsteady lid-driven cavity flow. *J. Fluid Mech.* **905**, A15.
- SCHLEGEL, M. & NOACK, B.R. 2015 On long-term boundedness of Galerkin models. *J. Fluid Mech.* **765**, 325–352.

- SCHMID, P.J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHMIDT, O.T. 2020 Bispectral mode decomposition of nonlinear flows. *Nonlinear Dyn.* **102** (4), 2479–2501.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structure. Part I, II, III. *Q. Appl. Maths* **XLV** (3), 561–590.
- SVANBERG, K. 2014 A class of globally convergent optimization methods based on conservative convex separable approximations. *SIAM J. Optim.* **12** (2), 555–573.
- SYMON, S., ILLINGWORTH, S.J. & MARUSIC, I. 2020 Energy transfer in turbulent channel flows and implications for resolvent modelling. *J. Fluid Mech.* **911**, A3.
- TERRAGNI, F., VALERO, E. & VEGA, J.M. 2011 Local POD plus Galerkin projection in the unsteady lid-driven cavity problem. *SIAM J. Sci. Comput.* **33** (6), 3538–3561.
- TIBSHIRANI, R. 1996 Regression shrinkage and selection via the LASSO. *J. R. Stat. Soc. B* **58** (1), 267–288.
- TIBSHIRANI, R.J. 2013 The LASSO problem and uniqueness. *Electron. J. Stat.* **7**, 1456–1490.
- TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.