

The l_1 -based sparsification of energy interactions in unsteady lid-driven cavity flow

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

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

(Received 10 April 2020; revised 6 July 2020; accepted 16 August 2020)

In this paper, sparsity-promoting regression techniques are employed to automatically identify from data relevant triadic interactions between modal structures in large Galerkin-based models of two-dimensional unsteady flows. The approach produces interpretable, sparsely connected models that reproduce the original dynamical behaviour at a much lower computational cost, as fewer triadic interactions need to be evaluated. The key feature of the approach is that dominant interactions are selected systematically from the solution of a convex optimisation problem, with a unique solution, and no *a priori* assumptions on the structure of scale interactions are required. We demonstrate this approach on models of two-dimensional lid-driven cavity flow at Reynolds number $Re = 2 \times 10^4$, where fluid motion is chaotic. To understand the role of the subspace utilised for the Galerkin projection in the sparsity characteristics, we consider two families of models obtained from two different modal decomposition techniques. The first uses energy-optimal proper orthogonal decomposition modes, while the second uses modes oscillating at a single frequency obtained from discrete Fourier transform of the flow snapshots. We show that, in both cases, and despite no *a priori* physical knowledge being incorporated into the approach, relevant interactions across the hierarchy of modes are identified in agreement with the expected picture of scale interactions in two-dimensional turbulence. Yet, substantial structural changes in the interaction pattern and a quantitatively different sparsity are observed. Finally, although not directly enforced in the procedure, the sparsified models have excellent long-term stability properties and correctly reproduce the spatio-temporal evolution of dominant flow structures in the cavity.

Key words: computational methods, low-dimensional models

1. Introduction

In the classical description of developed turbulent flows (Lumley 1979; Pope 2001; Jiménez 2018), energy is transferred across the hierarchy of coherent structures via nonlinear triadic interactions. Implicit in this picture is the fact that not all interactions have the same importance, but they occur in preferential patterns. In fact, extensive numerical evidence suggests that the nonlinear interaction pattern among coherent structures is sparse. 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

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

the influence of interactions with the complementary set of structures can be generally neglected with minor global effects.

Successful attempts to construct a reduced set of equations that exploit this sparsity have been made in the past, often for canonical geometries where triadic interactions are conveniently examined in Fourier space and using a coarse-grained partitioning of the hierarchy of scales. Laval, Dubrulle & Nazarenko (1999) considered two-dimensional homogeneous decaying turbulence and developed a reduced set of coupled partial differential equations governing the evolution of the large and small scales. In this model, only dominant terms were retained based on observations from direct numerical simulation. With the goal of identifying fundamental mechanisms underlying wall turbulence, Thomas *et al.* (2015) developed nonlinear reduced models of plane Couette flow directly from the governing equations by first partitioning the flow into a streamwise-averaged mean and a perturbation field, and then neglecting nonlinear interactions among the streamwise varying perturbations, i.e. the perturbation–perturbation nonlinearity (Thomas *et al.* 2014). The models captured well-established roll–streak dynamical features of wall turbulence and its statistics in a computationally efficient framework. The models also sustained turbulent dynamics down to minimal configurations where interactions between the streamwise mean flow and only one single streamwise wavenumber are retained.

When reduced-order dynamical representations are derived using Galerkin projection on a low-dimensional subspace identified by a set of modal structures (Fletcher 1984; Rowley & Dawson 2017), triadic interactions are conveniently studied in modal space by examining a third-order coefficient tensor arising from projection of the basis function on the convective term of the Navier–Stokes equations (Noack *et al.* 2008; Noack, Morzynski & Tadmor 2011). Sparsity characteristics have also been observed in this reduced-order setting. Couplet, Sagaut & Basdevant (2003) constructed Galerkin models of the separated turbulent flow past a backward-facing step using proper orthogonal decomposition (POD) modes (Lumley 1970; Sirovich 1987) and observed that the energy transfer pattern in modal space shares many properties with its counterpart in isotropic homogeneous three-dimensional turbulent flows (Yeung, Brasseur & Wang 1995). For instance, the authors observed that interactions are local in modal space and that a direct energy cascade exists. Analogously, Rempfer & Fasel (1994a) examined the power budget of POD modes in a transitional boundary layer and observed that interactions in modal space occur predominantly between triads of modes whose sum of modal indices is equal to zero, similar to energy interactions between Fourier modes in homogeneous turbulence. However, classical model order reduction techniques (Rowley & Dawson 2017) have not traditionally exploited this feature. In fact, when modal decompositions such as POD are employed, densely connected models are usually obtained, as the third-order coefficient tensor is dense (i.e. most coefficients are different from zero) for inhomogeneous flows without particular symmetries. This hinders the interpretation of the underlying physics of scale interactions and increases computational costs, since all triadic interactions have to be evaluated for advancing the model forward in time.

The first contribution of this work is that we propose to apply data-driven techniques (Blum & Langley 1997; Brunton, Proctor & Kutz 2016; Loiseau & Brunton 2018; Brunton, Noack & Koumoutsakos 2019) as a means to identify relevant triadic interactions in Galerkin models of turbulent flows. Our aim is to generate reduced-order models resolving a wide range of scales while preserving computational efficiency and interpretability by pruning weak interactions that are not relevant for the dynamics. The cornerstone of the proposed approach is l_1 -regularised regression (Friedman, Hastie & Tibshirani

2008; Tibshirani 2013), widely used in the statistical community to extract parsimonious representation of complex datasets containing a subset of predominant features. The non-differentiable, yet convex, nature of the l_1 regularisation allows transforming of the interaction selection problem into a convex optimisation problem that can be solved efficiently, with a unique solution. Since no *a priori* knowledge of the dynamics is utilised, the approach is fine grained and relevant interactions are identified in a mode-by-mode fashion across the hierarchy of modes. Sparsity-promoting regression techniques have been recently proposed by Brunton and coworkers (Brunton *et al.* 2016; Kaiser, Kutz & Brunton 2018) in the SINDy framework (sparse identification of nonlinear dynamics), as a means to discover parsimonious dynamical representations of systems whose underlying (but hidden) evolution equations are sparse in the space of possible functions (Brunton *et al.* 2016). Our work deviates from these efforts in the perspective. When formulated in partial differential form, the Navier–Stokes equations for incompressible flows are indeed structurally sparse, as only few terms – convection, viscous diffusion and pressure forces to conserve mass – participate in the overall equilibrium. However, when Galerkin models are derived, such structural sparsity is generally lost. What is preserved is the sparsity in the interaction pattern between scales of motion that emerges *a posteriori* in turbulent realisations. Fundamentally, we aim to exploit this feature and distil a structurally sparse model that reproduces the original behaviour. In addition, sparsity identification methods have been applied, so far, to relatively small Galerkin models (Loiseau & Brunton 2018), and it is not yet understood if these can be utilised to identify and extract relevant interactions in larger models in agreement with the established picture of energy interactions in turbulent flows. In this sense, our approach is closer to the recent work of Nair & Taira (2015), Taira, Nair & Brunton (2016) and Nair, Brunton & Taira (2018). These authors employed network-theoretic sparsification approaches (Newman 2018) to identify key vortex-to-vortex interactions in two-dimensional homogeneous turbulence, obtaining sparse models that capture the essential physics of unsteady fluid flow with a reduced number of interactions between the same large number of states.

The second contribution of this paper is that we examine how sparsity of energy interactions depends on the subspace used to generate the Galerkin model. Finding an appropriate subspace for projection is recognised as a challenging task (Noack *et al.* 2016), and several modal decompositions have been proposed differing in spirit and approach (see Taira *et al.* (2017) for a recent review). However, the role of the subspace on the organisation of energy interactions has not been explored in the past. To address this question, we examine and compare in this paper energy interactions and sparsity features of two families of Galerkin models. The first uses energy-optimal POD modes while the second uses modes oscillating temporally at a single frequency, obtained using a procedure based on spectral proper orthogonal decomposition (Sieber, Paschereit & Oberleithner 2016) and equivalent to a discrete Fourier transform (DFT) of the velocity snapshots. Here, we aim to understand if the optimal data-representation property of POD also provides the best description in terms of sparsity, even if POD is known to couple flow structures at different spatial or temporal scales (Noack *et al.* 2016; Towne, Schmidt & Colonius 2018).

This manuscript is organised as follows. For completeness, § 2 summarises the methodology utilised to generate reduced-order models using Galerkin projection, and then discusses how energy interactions in Galerkin models can be examined. Subsequently, the l_1 -based sparse regression approach is outlined and conceptual differences between our approach and the SINDy approach proposed in Brunton *et al.* (2016) are reported. In § 3, we demonstrate this methodology by considering relatively large Galerkin models of two-dimensional lid-driven cavity flow at a Reynolds number $Re = 2 \times 10^4$, where

the dynamics is chaotic (Auteri, Parolini & Quartapelle 2002). We first focus on modal decomposition of the flow and then move to energy analysis and sparsification. Conclusions are offered in § 4.

2. Methodology

2.1. Reduced-order modelling

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

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

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

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

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, with $\|\boldsymbol{\phi}_i(\mathbf{x})\| = 1$. These modes may be computed *a posteriori* from numerical or experimental data or *a priori* from a characteristic operator of the system (Taira *et al.* 2017) or from completeness considerations (Noack & Eckelmann 1994). Reduced-order models are then derived by projecting the governing equations onto the subspace defined by the modes (Rowley & Dawson 2017). Restricting our analysis to configurations where the boundaries are either no-slip walls or periodic, this procedure results in an autonomous system of coupled nonlinear ordinary differential equations

$$\sum_{j=1}^N M_{ij}\dot{a}_j(t) = \tilde{C}_i + \sum_{j=1}^N \tilde{L}_{ij}a_j(t) + \sum_{j=1}^N \sum_{k=1}^N \tilde{Q}_{ijk}a_j(t)a_k(t), \quad i = 1 \dots, N, \tag{2.3}$$

defining the temporal evolution of the coefficients $a_i(t)$. Here, we only report the definitions of the quadratic coefficients

$$\tilde{Q}_{ijk} = (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j \cdot \nabla \boldsymbol{\phi}_k), \tag{2.4}$$

while expressions for the tensors $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{L}}$ can be found in Noack *et al.* (2011). The matrix \mathbf{M} , with entries $M_{ij} = (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j)$, takes into account the fact that the spatial modes may not be orthogonal and is introduced here for generality.

If the N modes span collectively an N -dimensional subspace, M_{ij} is invertible and the system (2.3) can be rearranged as

$$\dot{a}_i(t) = C_i + \sum_{j=1}^N L_{ij}a_j(t) + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk}a_j(t)a_k(t) \quad i = 1 \dots, N, \tag{2.5}$$

with

$$C_i = \sum_{q=1}^N M_{iq}^{-1} \tilde{C}_q, \quad L_{ij} = \sum_{q=1}^N M_{iq}^{-1} \tilde{L}_{qj} \quad \text{and} \quad Q_{ijk} = \sum_{q=1}^N M_{iq}^{-1} \tilde{Q}_{qjk}. \quad (2.6a-c)$$

As observed by Rempfer & Fasel (1994a), the infinite-dimensional matrix M_{ij} should be first inverted and then truncated to maintain a good prediction accuracy. For the cases discussed in this paper, we have not followed this procedure as we observed that the matrix M_{ij} has a strong diagonal structure. Hence, the error performed by truncating it to size (N, N) and then inverting it can be reasonably assumed to be small.

Since the spatial modes satisfy automatically the boundary conditions, 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 $e_i(t) = \frac{1}{2} a_i(t) a_i(t)$, $i = 1, \dots, N$. The instantaneous rate of change is given by

$$\dot{e}_i(t) = C_i a_i(t) + \sum_{j=1}^N L_{ij} a_i(t) a_j(t) + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk} a_i(t) a_j(t) a_k(t), \quad i = 1, \dots, N, \quad (2.7)$$

obtained by multiplying (2.5) by $a_i(t)$. Note that, in a general case where the modes do not form an orthonormal set, the domain integral of the kinetic energy of velocity fluctuations is given by

$$E(t) = \frac{1}{2} \int_{\Omega} \mathbf{u}'(t, \mathbf{x})^2 \, d\Omega = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N M_{ij} a_i(t) a_j(t), \quad (2.8)$$

and not by a straightforward sum of the terms $e_i(t)$. The right-hand side of (2.7) is composed of three terms describing energy transfers between the hierarchy of modes. The first two describe variations of energy due to production/dissipation arising from interactions with the mean flow and from viscous effects (Noack *et al.* 2011). The third term defines variations of energy arising from inviscid nonlinear interactions between triads of modes. Following Rempfer & Fasel (1994a), these are defined in a time-averaged sense by the quadratic interaction tensor \mathbf{N} with entries

$$N_{ijk} = Q_{ijk} \overline{a_i a_j a_k}, \quad (2.9)$$

where the overbar denotes temporal averaging. The study of this term is the principal focus of the current analysis.

Spatial modes obtained from classical decompositions have generally global support over the domain (see e.g. Taira *et al.* 2017). The result is that the evolution equations (2.5) are not strictly sparse in the sense employed by Brunton *et al.* (2016). In fact, unless particular symmetries apply, the tensor \mathbf{Q} is generally dense, i.e. most of its entries are different from zero and the right-hand side of (2.5) contains all monomial terms in the modal amplitudes $a_i(t)$ up to order two.

However, as anticipated in the introduction, in turbulent realisations of the Navier–Stokes equations only a subset of all triadic interactions contributes to a significant degree to the overall energy budget (Rempfer & Fasel 1994b; Couplet *et al.* 2003). In this sense, sparsity is a primarily an *a posteriori* feature of solutions, i.e. a feature of the quadratic interaction tensor \mathbf{N} .

The approach developed in this work starts from this fundamental observation and aims to generate a sparse Galerkin model, defined by a sparse coefficient tensor \mathbf{Q}^s that is a good approximation of the original dynamical system in the sense that the mismatch between the transfer tensors \mathbf{N}^s and the original \mathbf{N} obtained from the definitions (2.4) and (2.6a–c) is as small as possible across the hierarchy of modes.

2.2. Sparse regression

To construct a sparse Galerkin system, we use a procedure akin to that utilised in previous work for calibrating Galerkin models from data (Perret, Collin & Delville 2006; Cordier, El Majd & Favier 2010; Xie *et al.* 2018) and more recently for the identification of sparse dynamical systems (Brunton *et al.* 2016). In the first step, we assume that N_t snapshots of the velocity field are available from simulation and arrange samples of the temporal coefficients $a_i(t_j)$, $i = 1, \dots, N$ and $j = 1, \dots, N_t$, into the data matrix $\mathbf{A} \in \mathfrak{R}^{N_t \times N}$, with entries $A_{ij} = a_i(t_j)$. Similarly, we construct the modal acceleration matrix $\dot{\mathbf{A}} \in \mathfrak{R}^{N_t \times N}$, containing the time derivative of the temporal coefficients obtained by projecting the modes $\phi_i(\mathbf{x})$ on snapshots of the Eulerian acceleration field $\partial_t \mathbf{u}(t_j, \mathbf{x})$ and correcting such projections with \mathbf{M} when modes are not orthogonal (see also Rempfer & Fasel 1994b). We then exploit the polynomial structure of the Galerkin system (2.5) to construct the database matrix $\Theta(\mathbf{A}) \in \mathfrak{R}^{N_t \times q}$

$$\Theta(\mathbf{A}) = \begin{pmatrix} 1 & a_1^1 & a_2^1 & \dots & a_N^1 & a_1^1 a_1^1 & \dots & a_N^1 a_N^1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & & \vdots \\ 1 & a_1^{N_t} & a_2^{N_t} & \dots & a_N^{N_t} & a_1^{N_t} a_1^{N_t} & \dots & a_N^{N_t} a_N^{N_t} \end{pmatrix}, \tag{2.10}$$

called the nonlinear feature library in Brunton *et al.* (2016), where $q = (N + 1) + N(N + 1)/2$ is the total number of features, the sum of constant, linear and quadratic interactions. The number of quadratic coefficients is only $N(N + 1)/2$ because the interaction between mode i and j is considered only once in (2.10). As discussed later on in the paper, this avoids columns of $\Theta(\mathbf{A})$ becoming linearly dependent, which would in turn result in numerical stability issues in the solution regression problem (see e.g. Perret *et al.* 2006; Cordier *et al.* 2010).

Arranging the projection coefficients tensors \mathbf{C} , \mathbf{L} and \mathbf{Q} associated with the i th mode into a coefficient vector $\beta_i \in \mathfrak{R}^q$, the Galerkin system (2.3) can be equivalently expressed as

$$\dot{\mathbf{A}}_i = \Theta(\mathbf{A})\beta_i, \quad i = 1, \dots, N, \tag{2.11}$$

where $\dot{\mathbf{A}}_i$ is the i th column of the modal acceleration matrix. The key idea is that if some nonlinear interactions are more important than others, then the corresponding entries of the coefficient vector β_i can be shrunk to zero with minor effects on the predictive ability of the resulting model. The challenge is to find a systematic method to identify the dominant interactions and prune unnecessary coefficients whilst calibrating the remaining model coefficients such as to preserve the overall energy budget. Here, we adopt an established sparsity-promoting regression technique known as LASSO regression (least absolute shrinkage selection operator, see Tibshirani 1996). In short, it leads to a set of N optimisation problems of the form

$$\min_{\beta_i} \|\Theta(\mathbf{A})\beta_i - \dot{\mathbf{A}}_i\|_2^2 + \gamma_i \|\beta_i\|_1, \quad i = 1, \dots, N, \tag{2.12}$$

one for each mode, where $\|\cdot\|_p$ denotes the l_p norm of a vector. The first term in the objective function in (2.12) produces calibrated models that have minimum prediction error on the modal acceleration (see discussion in Couplet, Basdevant & Sagaut 2005; Cordier *et al.* 2010). The second term penalises large model coefficients, regularises the regression and encourages sparsity in the solution by shrinking exactly to zero coefficients in β_i corresponding to columns of $\Theta(\mathbf{A})$ with little dynamical influence. Ideally, to prune unnecessary coefficients, a penalisation term proportional to the cardinality of β_i , $\text{card}(\beta_i)$, would formally be more correct (Jovanović, Schmid & Nichols 2014). However, the resulting optimisation problem would be computationally intractable even for Galerkin models of modest dimensions. In fact, this penalisation is usually relaxed to the computationally tractable l_1 term (Ramirez, Kreinovich & Argaez 2013). Regardless, the optimisation problems (2.12) are convex and thus have a unique solution. In addition, the approach lends naturally to parallelisation, since the optimisation problems can be solved independently for each mode. In initial stages of the research, we have found approaches based on sequential thresholded least squares (Brunton *et al.* 2016; Loiseau & Brunton 2018; Zhang & Schaeffer 2019) to be not sufficiently robust. Hence, solutions of (2.12) have been computed using the *sklearn* (Pedregosa & Varoquaux 2011) library, which implements a sub-gradient descent algorithm to manage the non-differentiability of the l_1 norm.

The weights γ_i in (2.12) are arbitrary and can be tuned to trade prediction ability (when they are small) for sparsity (when they are large). To formalise these concepts, we introduce the global reconstruction error ϵ

$$\epsilon = \sum_{i=1}^N \frac{\|\Theta(\mathbf{A})\beta_i - \dot{\mathbf{A}}_i\|_2^2}{\|\dot{\mathbf{A}}_i\|_2^2}, \quad (2.13)$$

and the system density ρ

$$\rho = \frac{1}{Nq} \sum_{i=1}^N \text{card}(\beta_i). \quad (2.14)$$

In (2.13), the absolute reconstruction error $\|\Theta(\mathbf{A})\beta_i - \dot{\mathbf{A}}_i\|_2^2$ is normalised with the mean squared acceleration $\|\dot{\mathbf{A}}_i\|_2^2$ to balance the global reconstruction error across the hierarchy, which would be otherwise dominated by the most energetic modes. On the other hand, the density ρ ranges from 0, when all interactions have been pruned, to 1, for a fully connected model. Note that for large models, the density is dominated by the quadratic tensor \mathbf{Q} . A one-parameter family of models can be generated by varying the regularisation weights γ_i , producing a Pareto front (Schmidt & Lipson 2009) on the ρ - ϵ plane. Since only a subset of triadic interactions is relevant, the expectation is that a sweet spot appears on this curve, defining ‘optimal’ penalisation coefficients γ_i . It is important to observe that the penalisation coefficient γ_i can be chosen independently for each index i , implying that reconstruction error and sparsity can be modulated arbitrarily across the spectrum of modes. In our analysis we consider two different modulation strategies. In strategy S1, the weight is constant for all modes, $\gamma_i = \gamma$. This strategy sparsifies more aggressively the equations of motion of low-energy modes, because the l_1 penalisation term has a higher importance than the l_2 component. In this work we observed that the lowest global reconstruction error is obtained when γ_i is kept constant across the modes. We also introduce strategy S2, where the weight is normalised with respect to the mean squared modal acceleration as $\gamma_i = \|\dot{\mathbf{A}}_i\|_2^2 \gamma$. This is equivalent to solving problem (2.12) using the relative error in (2.13) as least-squares component of the objective function. This strategy

results in a more balanced sparsification across the hierarchy of modes and avoids earlier truncation, i.e. when all coefficients of a high-index mode are set to zero. Other strategies can be, of course, devised. Here, we mention, for instance, the possibility to tune the penalisation coefficients to obtain a uniform sparsification across the spectrum or to obtain a uniform relative reconstruction error. Analysing these strategies is an interesting avenue for future work.

One potential modification of this approach is that discussed in Loiseau & Brunton (2018), namely to enforce that the nonlinear term in the sparsified Galerkin model conserves energy exactly (see e.g. Balajewicz, Dowell & Noack (2013) and Noack *et al.* (2008) for a formal definition). In practice, this can be achieved by introducing a set of constraints on the coefficients vectors β_i . The constraints, however, couple together the regression problems of all modes, resulting in one optimisation problem of larger dimension. As we will demonstrate later in the paper, the energy conservation error of models obtained from the unconstrained approach is small in relative terms. This occurs because the temporal coefficients in \mathbf{A} are originally obtained from an energy conserving nonlinearity, and the regression ‘discovers’ this property from data. Hence, throughout this paper we always solved problems (2.12) independently, without additional constraints.

3. Results

We now apply this methodology to two-dimensional unsteady flow in a lid-driven square cavity. This is 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; Arbabi & Mezić 2017; Fick *et al.* 2018), and we thus consider it here as an exemplar to demonstrate the ideas discussed in the introduction.

The Reynolds number is defined as $Re = LU/\nu$ where L , U are the cavity dimension and the lid velocity, respectively, while ν is the kinematic viscosity. These quantities are used to make the equations of motion non-dimensional. We purposefully investigate a regime at $Re = 2 \times 10^4$, where the flow evolves in a chaotic manner (Auteri *et al.* 2002; Peng, Shiau & Hwang 2003, see also the animation of the vorticity field in the supplementary movie is available at <https://doi.org/10.1017/jfm.2020.707>). The chaotic nature of the problem ensures that the frequency spectrum of velocity fluctuations is continuous and energy transfers are scattered in modal space, rather than being highly organised as for periodic flows (Noack *et al.* 2011). 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 $\omega = \partial v / \partial x - \partial u / \partial y$.

Direct numerical simulations (DNS) were performed in OpenFOAM with the incompressible flow solver icofoam. The convective and viscous terms are spatially discretised with a second-order finite-volume technique and the temporal term with a semi-implicit Crank–Nicholson scheme. Special treatments of the singularities developing at the top corners due to the discontinuity in the velocity boundary conditions (Botella & Peyret 1998) were not deemed necessary. A grid independence study was initially 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, the high vorticity filaments characteristic of two-dimensional turbulence as well as the spatial structure of the lowest-energy modes utilised for the projection. Similar grid resolutions have been used by Cazemier *et al.* (1998) at similar Reynolds numbers.




FIGURE 1. Vorticity field ω of three different snapshots separated by one non-dimensional time unit, increasing from left to right. An animation of the vorticity field, illustrating the chaotic nature of the dynamics and the dominant flow features is available as a supplementary movie.

Three snapshots of the vorticity field obtained from these simulations are shown in figure 1 (see also the animation in the supplementary movie). 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. The characteristic non-dimensional frequency of this wave-like motion is $f = 0.7$. From simulation, we extract $N_t = 1500$ velocity snapshots using a non-dimensional sampling period $\Delta t = 0.1$. These settings are sufficient to adequately time resolve the fast scales as well as to include many shedding events at the bottom-right corner, making the regression problems (2.12) statistically reliable.

3.1. Modal decomposition

First, we consider models generated using POD modes. POD produces economic reduced-order models, but has the well-known shortcoming of mixing together fluid motions at different temporal/spatial scales (Mendez, Balabane & Buchlin 2019). Second, we consider models generated from modes oscillating at a single frequency obtained from a procedure that is equivalent to a DFT of the velocity snapshots. For practical convenience, we obtain the two distinct sets of modes using the same computational technique, based on the approach proposed by Sieber *et al.* (2016) which only operates on the temporal correlation matrix. Briefly, the method considers the temporal correlation matrix $\mathbf{R} \in \mathbb{R}^{N_t, N_t}$, with entries

$$R_{ij} = \frac{1}{N_t} (\mathbf{u}'(t_i, \mathbf{x}), \mathbf{u}'(t_j, \mathbf{x})), \tag{3.1}$$

and then defines a filtered correlation matrix \mathbf{S} , with elements

$$S_{ij} = \sum_{k=-N_f}^{k=N_f} g_k R_{i+k, j+k}, \tag{3.2}$$

given by the application of the filter coefficient vector \mathbf{g} along the diagonals of the correlation matrix. An ordered set of temporal coefficients $\mathbf{a}_i = [a_i(t_1), \dots, a_i(t_{N_t})]$ and

associated mode energies λ_i is then obtained from the eigendecomposition of \mathbf{S} ,

$$\mathbf{S}\mathbf{a}_i = \lambda_i\mathbf{a}_i, \quad (3.3)$$

so that $\lambda_i\delta_{ij} = \mathbf{a}_i^\top \cdot \mathbf{a}_j$. As discussed in Sieber *et al.* (2016), when the filter is extended over the entire dataset and in the limit of number of samples tending to infinity, the filtered correlation matrix converges to a Toeplitz, circulant matrix. Then, its eigenvalues trace the power spectral density of the underlying dataset. On the other hand, the eigenvectors \mathbf{a}_i corresponds to the Fourier basis. This procedure generates conjugate pairs of modal structures with same energy oscillating at a single frequency. These can be viewed as a set of modal oscillators exhibiting periodic fluctuations (Taira *et al.* 2017) and tracing fluid motion at on a two-dimensional subspace. In practice, for a finite-length dataset, we filter the temporal correlation matrix assuming periodicity using a box-car filter, as suggested in Sieber *et al.* (2016). Hereafter, we will refer to the modal structures identified by this procedure as DFT modes.

One important consideration is that, unlike dynamic mode decomposition (see Rowley *et al.* 2009; Schmid 2010; Chen, Tu & Rowley 2012), DFT lacks the ability to discern and identify dominant frequency components. Instead, a number of modes equal to the number of snapshots utilised is produced, oscillating in conjugate pairs at specific frequencies determined by the sampling period Δt and observation time T (Mendez *et al.* 2019). This property, picket fencing, results in frequencies that are integer multiples of the fundamental frequency $f_1 = T^{-1}$, up to the Nyquist component $f_{Nyq} = (2\Delta t)^{-1}$. In addition, unlike for POD, as the length of the dataset is increased, the number of energy-relevant modes increases and low-frequency modes with little dynamical importance appear. The approach we use here is to divide the dataset into five partition of thirty time units, covering an average of 20 cycles of the dominant oscillatory component, and providing sufficient frequency resolution to distinguish small scale spectral features. In addition, two possible ways of sorting pairs of modal structures are possible, i.e. by energy content (using the eigenvalues λ_i) or by frequency. Models obtained with the two sorting schemes will be referred to as DFT_e and DFT_f , respectively.

We now focus on the characteristics of the modal structures obtained by these two methods. We denote the normalised cumulative sum of the eigenvalues λ_i of the (filtered) correlation matrix as

$$e(n) = \sum_{i=1}^n \lambda_i / \sum_{i=1}^{N_f} \lambda_i, \quad (3.4)$$

describing the fraction of the fluctuation kinetic energy captured by the first n elements of the expansion (2.2).

This quantity is shown in figure 2(a) for the POD and for the two possible DFT sorting schemes. As expected, a larger energy is captured by the POD basis. For the DFT decomposition, the energy-based sorting is more efficient at data compression, although the difference vanishes for large n , since for low-energy modes the two sorting schemes are equivalent. The modal energies associated with the DFT_f modes are shown in figure 2(b) as a function of the modal index i . The distribution is characterised by a continuous component, with modal energy decaying with frequency, and a discrete component, with a fundamental peak for the pair of modes (31, 32) and its first few harmonics. The peak, at a non-dimensional frequency $f = 0.7$, is physically originated from the high-energy structures transported along the shear layer by the rotation of the main vortical structure.

This can be observed in figures 3(a) and 3(b), showing the vorticity field ω of the DFT mode pair (31, 32). This pair of modes describes a vorticity perturbation having




FIGURE 2. (a) Cumulative sum of the first 100 eigenvalues of S_{ij} for the three decompositions considered. (b) Distribution of the modal energies of DFT modes sorted by frequency.

the form of a wave travelling along the edge of the main vortex. Hence, the spatial structure of the two modes is shifted in the direction of the shear layer by half-wave. Travelling-wave structures in cavity flows have already been observed in simulation by Poliashenko & Aidun (1995); Auteri *et al.* (2002) and characterised by global stability analysis and Koopman analysis by Boppana & Gajjar (2010) and Arbabi & Mezić (2017), respectively. The two leading POD modes, reported in figure 3(c,d), have the same energy and capture the same travelling-wave pattern described by the leading DFT mode pair.

3.2. Energy analysis

To provide a more robust foundation to understand the sparsification results reported in §§ 3.3 and 3.6, we first focus on the analysis of the average energy interactions. The structure of the interaction tensor \mathbf{N} for a large POD-based model with $N = 75$, reconstructing more than 99% of the fluctuation kinetic energy, is reported in figure 4, showing the magnitude of the interactions for three slices for $i = 1, 10$ and 75 , in panels (a), (b) and (c), respectively. All entries of the tensor \mathbf{N} are generally non-zero, although the strength of the interactions varies across several orders of magnitude. This is a combined result of the projection coefficients tensor \mathbf{Q} (shown later), whose entries are typically non-zero, and of the complex spectral structure of the temporal coefficients $a_i(t)$. The most important feature of figure 4 is that interactions are highly organised and there exists a subset of interactions that are more active. Specifically, for any mode i , triadic interactions can be classified as illustrated in panel (a) in four different categories by introducing a cutoff modal index n . The subset of interactions denoted as LL corresponds to nonlinear energy transfer involving pairs of low-index modes, HL and LH denote interactions involving high–low/low–high index modes, while HH denotes the subset of interactions involving pairs of high-index modes. We observe that the areas corresponding to LL and HL/LH are the most active. If we map low/high modal indices to large/small scales, this result is in agreement with the picture of energy transfer 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. In addition, interactions are not symmetric with respect to a swap of indices j, k . This can be quantified




FIGURE 3. Vorticity field of the most energetic pair of DFT modes, panels (a,b), and of the first two POD modes, panels (c,d).

by computing the coefficient

$$\chi_i(n) = \frac{\sum_{j=1}^n \sum_{k=1}^N N_{ijk}}{\sum_{j=1}^N \sum_{k=1}^n N_{ijk}}, \tag{3.5}$$

representing the relative dynamical importance of the subset of interactions LL + HL and LL + LH. Figure 4(d) shows χ_i for $i = 1, 10$ and 75 as a function of the normalised cutoff n . The interaction subset HL is up to four times more important than the subset LH. This is a consequence of the asymmetry of the projection coefficients Q_{ijk} , which arise from the fact that the convective transport of structure $\phi_k(\mathbf{x})$ operated by the structure $\phi_j(\mathbf{x})$ is more intense when the modal structure $\phi_j(\mathbf{x})$ describes large-scale flow features.

We now consider energy analysis of a large, full-resolution DFT_f model constructed from five partitions of thirty time units as discussed in § 3.1. The model is composed of all $N = 300$ modes, corresponding to 150 distinct frequencies. We perform modal decomposition and energy analysis on each partition separately, and then average the mean energy transfer rate tensor \mathbf{N} over the five partitions. Figure 5(a) shows the mean transfer rate distribution for mode $i = 100$. Energy interactions in the DFT model are very sparsely distributed on a thin horseshoe-shaped structure composed of three branches (denoted in the figure as L , C and U) of 2×2 blocks, and all other mean energy transfer rates interactions are identically zero. This pattern results from the joint effect of the oscillatory nature of the temporal coefficients and the quadratic nonlinearity of system (2.3), which




FIGURE 4. Magnitude of the average interaction tensor coefficients N_{ijk} for three POD modes across the spectrum, $i = 1, 10$ and 75 in panels (a), (b) and (c) respectively, for a model resolving 99 % of the fluctuation kinetic energy. Panel (d) shows the coefficient $\chi_i(n)$ as a function of the normalised cutoff n for the same three modes.

can only be satisfied by triads of modes having matching temporal wavenumbers. A less pronounced horseshoe-shaped distribution of the energy interactions has been previously observed in energy analysis of POD-based models of three-dimensional transitional boundary layers Rempfer & Fasel (1994a,b). These authors noticed that low-energy modal structures resemble Fourier modes in the spanwise direction (justified by the spanwise periodic domain) and thus coefficients Q_{ijk} and energy interactions are non-zero only for specific triads of modes. In the present case, this pattern is determined exclusively by the temporal coefficients as the tensor \mathbf{Q} constructed from projection modes does not possess any sparsity structure and its coefficients have a similar statistical distribution to that obtained using the POD modes. This is illustrated in figure 6 showing maps of the first slice of the tensor \mathbf{Q} of the largest Galerkin models considered here, constructed from the POD and the DFT decompositions, in panels (a) and (b) respectively. We observe that no underlying structure is present except for the asymmetry already observed in the energy analysis in figure 4. This property is confirmed in the probability distribution of the coefficients, shown in panel (c).

To facilitate the interpretation of the energy interaction pattern, we follow Rempfer & Fasel (1994b) and Arbabi & Mezić (2017) and define oscillatory modal structures

$$\mathbf{u}_l(t, \mathbf{x}) = a_{2l-1}(t)\boldsymbol{\phi}_{2l-1}(\mathbf{x}) + a_{2l}(t)\boldsymbol{\phi}_{2l}(\mathbf{x}), \quad (3.6)$$

numbered by the index l and tracing fluid motion at a single frequency on a two-dimensional subspace. Their modal energy is

$$e_l(t) = \frac{1}{2}(a_{2l-1}^2(t) + a_{2l}^2(t)) + a_{2l-1}(t)a_{2l}(t)(\boldsymbol{\phi}_{2l-1}(\mathbf{x}), \boldsymbol{\phi}_{2l}(\mathbf{x})). \quad (3.7)$$

Numerical experiments show that, for large number of snapshots, pairs of modes $\boldsymbol{\phi}_{2l-1}(\mathbf{x})$ and $\boldsymbol{\phi}_{2l}(\mathbf{x})$ tend to be orthogonal. Hence, considering the evolution equation for the modal energy $e_l(t) \sim \frac{1}{2}(a_{2l-1}^2(t) + a_{2l}^2(t))$ leads to the condensed triadic interaction tensor $\hat{\mathbf{N}}$ of size $(N/2, N/2, N/2)$ with entries

$$\hat{N}_{lmn} = \sum_{i=l}^{l+1} \sum_{j=m}^{m+1} \sum_{k=n}^{n+1} N_{ijk}, \quad (3.8)$$

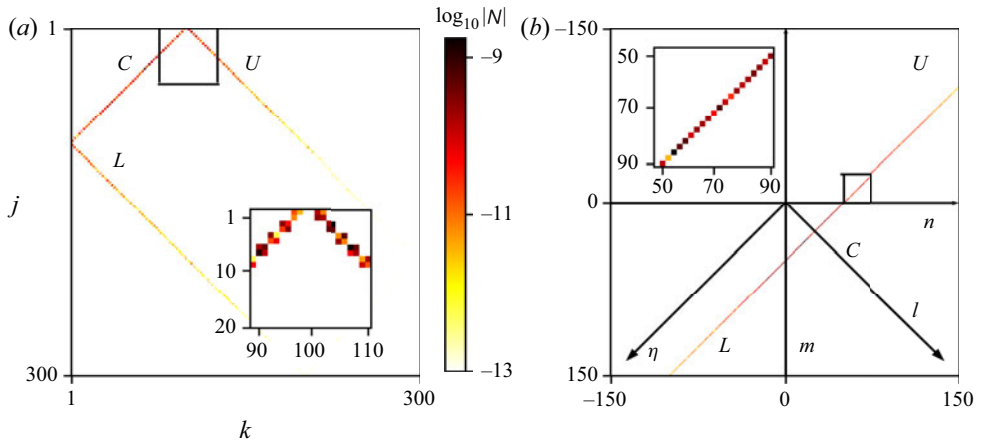


FIGURE 5. (a) Magnitude of the average interaction tensor N_{ijk} for $i = 100$, with the three characteristics branches, showing that active interactions come in 2×2 blocks corresponding to matching triads of modes. The small inset focuses on the interactions of branches C and U . (b) Magnitude of the average interaction tensor (3.8) where the three branches of panel (a) have been unfolded on a larger plane spanned by the coordinates l and η . The inset shows details of the interactions of the branch U in the plane $\eta - l$.

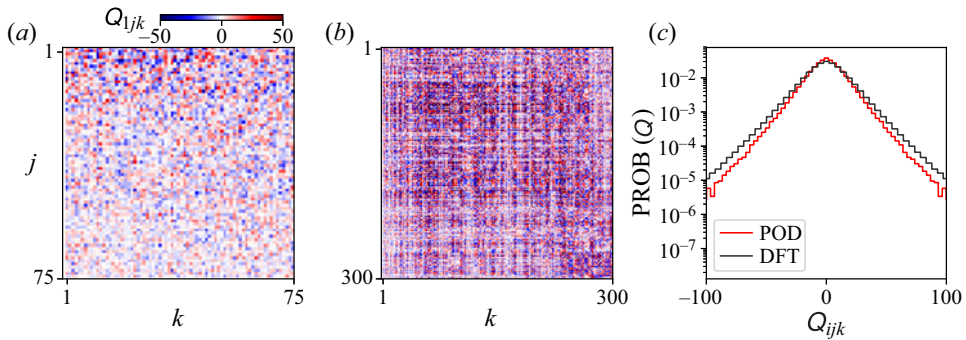


FIGURE 6. Maps of Q_{1jk} for Galerkin models constructed from the POD and the DFT_f decompositions, in panels (a) and (b), respectively. Panel (c) shows the probability distribution of all quadratic coefficients for these two models.

lumping together the 2×2 blocks of interactions at matching triads of figure 5(a). In addition, the three branches L , C and U can be unfolded and conveniently visualised on a two-dimensional plane spanned by the coordinate l , the modal structure index, and $\eta = m - n$, representing the distance in modal space between pairs of temporal wavenumbers. This unfolding process is shown in panel (b) of figure 5, and when repeated for all modal structures leads to the distribution shown in figure 7(a). In figure 7(b), we report the average transfer rate \hat{N}_{lmn} normalised with the total average transfer rate for each structure, the quantity

$$\hat{T}_l = \sum_{m=1}^{N/2} \sum_{n=1}^{N/2} \hat{N}_{lmn}, \quad l = 1, \dots, N/2, \tag{3.9}$$

to illustrate more clearly the relative strength of the interactions. In figure 7(c), the normalised mean transfer rate for $l = 1$ and 150 is reported. Interactions between

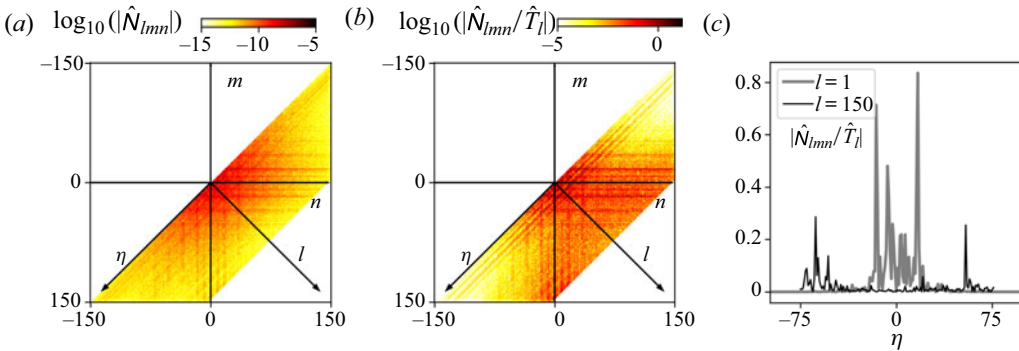


FIGURE 7. (a,b) Absolute and relative strength of the energy interactions between pairs of DFT modes for a model with $N = 300$ visualised on the plane m, n , with the additional coordinates l and η . (c) Relative energy interactions for the first and last mode pairs.

triads of pairs of DFT modes are organised in agreement with the physics of scale interactions previously discussed for POD models. In absolute terms, the most relevant interactions are clearly those located near the origin of the coordinates. These correspond to low-index modes where nonlinear interactions with other low-index modes dominate, while interactions with the high-index modes, for larger η , are less important. This suggests that a sparsification approach based on pruning the interactions involving the high-index modes, i.e. the small scales, would be effective. By contrast, for high-index modes, relevant energy interactions are organised in bands along the axes m and n and involve energy exchange between low-index modes and high-energy, high-index modes. This suggests that the dynamics of the small scales is driven primarily by non-local interactions with the largest structures of the flow and not by small-scale/small-scale quadratic interactions. The slight asymmetry visible in panel (c) arises from the structure of the coefficients tensor Q_{ijk} and has the same physical origin as that observed in figure 4 for the POD model.

3.3. Sparsification of POD-based models

We now apply the methodology presented in § 2 to three POD-based models resolving 90 %, 95 % and 99 % of the kinetic energy, respectively (see table 1 for details). Because the size of the database matrix $\Theta(\mathbf{A})$ grows quadratically with the number of modes, the number of possible interactions q can easily become larger than the number of available snapshots N_t , resulting in an underdetermined regression problem and overfitting. This is a well-understood issue in data analysis and requires cross-validation techniques to ensure the statistical reliability of the result (Friedman *et al.* 2008). In this work, we employed K -fold cross-validation, using typically $K = 10$. Briefly, the database is first divided into K folds. The model is trained using $K - 1$ blocks and the reconstruction error ϵ of (2.13) is obtained from the fold that was left out. This procedure is iterated over all folds, obtaining the mean and the standard deviation of ϵ .

Figures 8(a)–8(c) show the sparsification curves on the ρ – ϵ plane for the three POD models considered. The mean of ϵ across the folds is displayed as a thick black line, while the grey dashed line indicates plus or minus one standard deviation. These curves have been obtained by solving problem (2.12) using strategy S1 and progressively increasing the regularisation weight. When low weights are used, dense systems with good prediction accuracy are obtained. The opposite is true for large weights, identifying points in the left

| n | 1 | 5 | 10 | 15 | 20 | 26 | 50 | 75 | 80 | 95 | 100 | 300 |
|-------------------------------|------|------|------|------|------|------|------|------|-------|-------|-------|-----|
| POD | 0.26 | 0.74 | 0.85 | 0.89 | 0.9 | 0.95 | 0.98 | 0.99 | 0.995 | 0.998 | 0.999 | — |
| DFT _{e} | 0.17 | 0.49 | 0.62 | 0.74 | 0.8 | 0.88 | — | 0.92 | 0.95 | 0.97 | 0.98 | 1 |
| DFT _{f} | 0.02 | 0.07 | 0.15 | 0.39 | 0.43 | 0.47 | 0.91 | — | 0.96 | 0.97 | 0.98 | 1 |

TABLE 1. Normalised cumulative energy distribution $e(n)$ for POD and DFT modes, where the latter are sorted by energy content (DFT _{e}) or by frequency (DFT _{f}).




FIGURE 8. The ρ - ϵ curves for three POD models resolving 90 %, 95 % and 99 % of the kinetic energy in panels (a), (b) and (c), respectively. The black line represents the cross-validated error averaged over $K = 10$ folds. The dashed grey lines represent plus/minus one standard deviation of the cross-validated error calculated over the folds. The red dashed line shows the reconstruction error obtained with the greedy approach. The squares indicate the global reconstruction error of the Galerkin model obtained directly from projection.

part of the graph characterised by low density and poor prediction accuracy. As postulated in § 2, the curves show a sweet spot at around $\rho \approx 0.2$, displaying a plateau for $\rho \gtrsim 0.2$, while the error ϵ grows quickly when $\rho \lesssim 0.2$. These results indicate that it is possible to prune ~ 80 % of the quadratic interactions in model (2.3) without influencing the average prediction accuracy. The red dashed line represents the reconstruction error obtained with a naive sparsification approach. The approach consists in pruning coefficients of Q_{ijk} in the area denoted as HH in figure 4(a). The approach exploits the structure of N_{ijk} and is therefore referred to as ‘greedy’. By varying $n \in (1, N)$, models with different sparsity are obtained, with $n = N$ corresponding to the original projection model projection (indicated as a red square in figure 8). The shapes of the sparsification curves for the greedy approach are similar to those obtained with the l_1 regression. This is a direct consequence of the existence of a subset of most relevant energy interactions. However, the reconstruction error obtained from the greedy method is generally higher than that obtained by the l_1 regression, since the optimisation procedure involved in the l_1 approach modulates the strength of the remaining interactions by tuning the active quadratic coefficients, minimising the prediction error. As we show later in § 3.5 dedicated to analysing the model performance in time integration, this difference will have a marked effect on the long-term temporal stability of the models.

Regardless of the approach, the mean reconstruction error decreases as the resolved energy increases, moving from panel (a) to panel (c), as more modes participate in

capturing the dynamics of velocity fluctuations. In addition, larger models can be more effectively sparsified, as the sparsification curve drops more rapidly. This results from the non-local structure of energy interactions shown in figure 4. When one additional low-energy mode is included, the number of relevant interactions to be retained in the model only grows as $O(N)$ and not as $O(N^2)$, i.e. all non-local interactions with the rest of the hierarchy denoted as LL, LH and HL in figure 4(a). Since the total number of possible interactions grows as $O(N^3)$, larger models can be more effectively sparsified. This is conceptually in agreement with the observations of Taira *et al.* (2016) on the sparsification properties of discrete vortex models. We also observe that the mean prediction error does not necessarily decrease monotonically when the density increases. This phenomenon is particularly visible for the model in panel (b) but all models reproduce the same behaviour. This is a symptom that the number of available snapshots (1500) is potentially not large enough for the number of coefficients ($q = N \times (N + 1)/2 + N + 1 = 2926$ for the model in panel (c)) and overfitting would have occurred if no cross-validation had been performed.

3.4. Energy interactions identified by the regression and conservation properties

To visualise the sparsity pattern identified by the regression as the regularisation weight in (2.12) is increased (constant for all modes in strategy S1), we introduce the tensor $\boldsymbol{\gamma}$ with entries γ_{ijk} defined as the value of the regularisation weight at which the corresponding coefficient Q_{ijk} is shrunk to zero by the LASSO. Figures 9(a)–9(c) show three slices of $\boldsymbol{\gamma}$ for modes $i = 1, 10$ and 75 , respectively, for the largest POD model considered, capturing 99% of the total fluctuation kinetic energy. The first interactions to disappear are the small-scale/small-scale interactions. Increasing the penalisation, interactions that are local in modal space are progressively pruned, leaving only non-local interactions involving triadic exchanges with the low-index modes for large penalisations. Interestingly, this pattern does not change qualitatively nor quantitatively as the modal index i increases. In fact, a comparable number of interactions is retained across the hierarchy and the governing equations of all modes are sparsified by an equal amount. Hence, sparsification has not produced mode truncation, which would have occurred if all coefficients of some low-energy modes had been shrunk to zero by the LASSO. This behaviour can be justified by noting that the mean square acceleration $\|\mathbf{A}_i\|_2^2$ of the POD amplitude coefficients varies weakly with i . In fact, the sparsification pattern does not change significantly when strategy S2 is used. In figure 10(a–c) the base ten logarithm of the mean energy interaction tensor N_{ijk}^s computed as in (2.9) with the sparse coefficient tensor \mathbf{Q}^s is shown for the same three modal indices as in figure 9. The data refer to sparse models with $\rho = 0.3$, located nearby the sweet spot of the curves in figure 8(c). It can be observed that the sparsified model has a pattern of interactions resembling that of the dense model in figure 4. However, weak interactions and the associated flow physics have been pruned. It is also clear that the asymmetry of the interaction pattern observed in figure 4 and the physical mechanism that originates it are invisible to the regression and the interaction pattern in figure 10 is now symmetric with respect to a swap of the indices j, k .

Despite the aggressive pruning, the sparse models reproduce fairly accurately the overall structure of the intermodal energy budgets. In the present flow configuration, the convective nonlinearity is energy conserving and Galerkin models should obey the relation $\sum_{i=1}^N T_i = 0$ as $N \rightarrow \infty$, with $T_i = \sum_{j=1}^N \sum_{k=1}^N N_{ijk}$ the time-averaged energy transfer rate to/from mode i . For finite-dimensional approximations, this property is not satisfied exactly and the residual of the summation can be taken as a measure of the overall energy balance. Figure 11(a) shows such residual for the l_1 sparsified models (empty circles) and




FIGURE 9. Distribution of the base ten logarithm of γ_{ijk} for $i = 1, 10$ and 75 , in panels (a), (b) and (c), respectively, for the POD model resolving 99% of the fluctuation kinetic energy.

for the models obtained with the greedy approach (empty squares), as a function of the density. The residual is normalised by the root mean square value of the rate of change of the integral fluctuation energy. Note that the greedy model at $\rho = 1$ is the model obtained directly from projection. It can be observed that the energy conservation error is relatively small, of the order of $10^{-3} \div 10^{-4}$. For large densities, it is larger than that of the projection model, because the regression tunes model coefficients to minimise the mean square error on the modal accelerations and does not enforce this physical constraint directly. The energy conservation residual decreases for sparser models and is ten times smaller than the projection model, owing to the lower number of active coefficients that participate in the regression. Figure 11(b) shows the distribution of the time-averaged energy transfer rate associated with mode i for the l_1 sparsified model at $\rho = 0.3$ (red crosses) and the dense model obtained from projection (empty circles). Data are reported every second mode. For the projection model, the net energy transfer is negative for the first few modes and changes sign at $i \sim 10$. Physically, this trend suggests that the first few modes extract energy from the mean flow and feed the dissipative high-index modes via triadic interactions. The l_1 sparsified model correctly reproduces this global trend, even though no constraints have been introduced (Loiseau & Brunton 2018). As discussed in § 2, it is argued that this is a general properties of data-driven techniques relying on optimisation ideas, such as the LASSO, which naturally reproduce invariants and conservation properties embedded in the data to a level defined by noise levels. For instance, Taira *et al.* (2016) used network-theoretic ideas to sparsify connections in a discrete vortex model and observed that sparsification conserves the invariants of discrete vortex dynamics.

3.5. Long-term temporal behaviour of the l_1 sparse systems

We now turn our attention to the long-time behaviour of the sparsified models under temporal integration. We consider results for models resolving 95% of the turbulent kinetic energy as an illustrative example, and use the projections of the POD modes onto one of the DNS snapshots to obtain initial conditions. Results are shown in figure 12. Panel (a) shows the temporal evolution of the integral fluctuation kinetic energy $E(t)$, as defined in (2.7), for the dense projection model and the sparse models obtained with the l_1 and greedy approach, with same density $\rho = 0.3$. The temporal evolution obtained with these models is compared against the fluctuation kinetic energy from DNS. For the projection model, the integral kinetic energy grows substantially in the




FIGURE 10. Base ten logarithm of the sparsified interaction tensor N_{ijk}^s for $i = 1, 10$ and 75 in panels (a), (b) and (c), respectively, for a POD model resolving 99% of the fluctuation kinetic energy and $\rho = 0.3$.

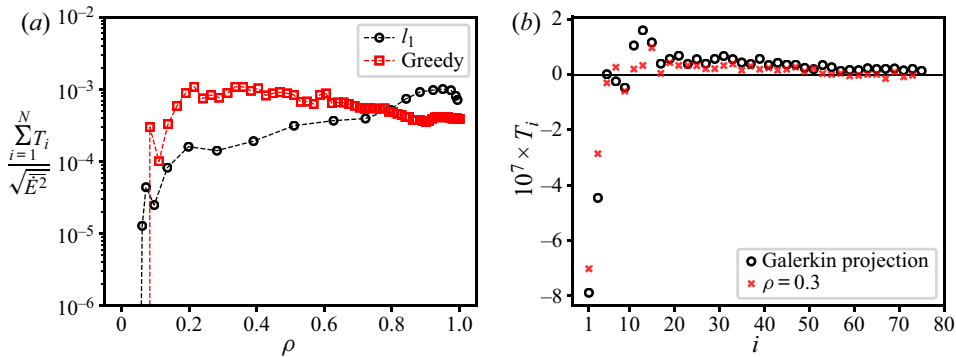


FIGURE 11. (a) The normalised energy conservation error as a function of the density ρ , for the greedy sparsification approach (empty squares) and the l_1 based approach (empty circles). (b) The total net energy transfer rate T_i as a function of the modal index i for two POD models resolving 99% of the kinetic energy, with coefficients identified from projection and for a l_1 sparse model. One every two data points are reported.

first 40 time units and then saturates on a value that is approximately two orders of magnitude larger than what observed in DNS. The over-prediction occurs because the truncation of the small scales in the ansatz (2.2) leads to a significant imbalance of the production–dissipation budget within the model (Noack, Papas & Monkewitz 2005; Noack *et al.* 2008, 2011; Balajewicz *et al.* 2013). A qualitatively similar behaviour, if not worse, is then necessarily observed for the model sparsified with the greedy approach, since neglecting weak interactions alone does not cure the original dissipation problems. Conversely, the l_1 model is able to predict the correct average fluctuation kinetic energy and has excellent long-term stability properties, despite this being not enforced in the regression procedure (see Fick *et al.* 2018). We argue that this is due to the fact that the l_1 procedure performs a ‘prune-then-calibrate’ approach, where weak interactions are first pruned and the remaining active coefficients are then tuned in the optimisation involved in (2.12) to match the reference dynamics. It is evident from these results that this second step is key to obtain accurate long-term behaviour. Figures 12(b)–12(d) show a shorter segment of state space trajectory from these models projected onto three different pairs

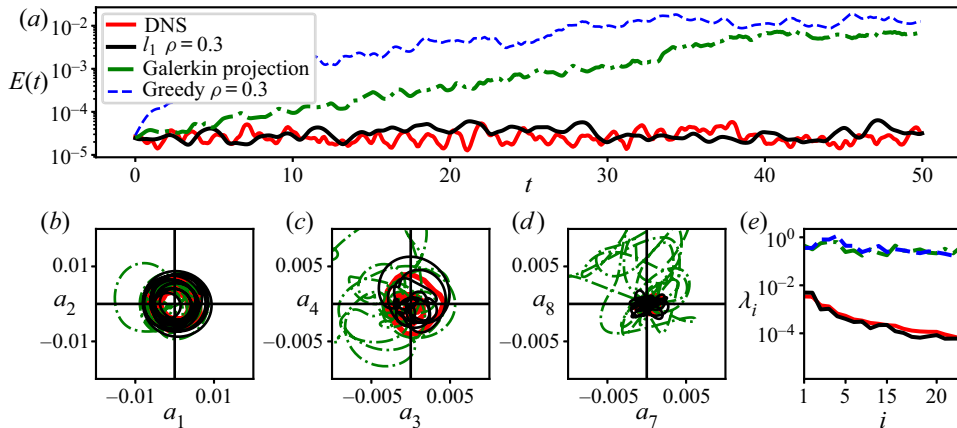


FIGURE 12. (a) Temporal evolution of the turbulent kinetic energy $E(t)$ from DNS compared to that obtained from temporal integration of the l_1 , greedy and Galerkin projection models. Sparse models have $\rho = 0.3$. (b)–(d) State space projections onto three different mode pairs. Data for the greedy model are omitted as the trajectory quickly leaves the visible range. (e) Average modal energy ($\lambda_i = \overline{a_i a_i}$) from DNS and long-time integration of the models considered in panel (a).

of modes. We omit the data for greedy approach since orbits quickly drift out of the visible range and this case is thus qualitatively similar to that from the dense model. It is clear that the trajectory of the l_1 model remains in the same volume of state space occupied by the DNS projections, while the projection model drifts away to a different region of state space, over-predicting the integral kinetic energy. This is more effectively visualised in panel (e), displaying the average modal energy $\lambda_i = \overline{a_i a_i}$ as a function of the modal index. Data are obtained from averaging long trajectories after the initial transient in panel (a) has completed. The projection and greedy models predict much larger energies across the entire spectrum, while the l_1 correctly predicts the correct decay of modal energies. It is, of course, not possible to guarantee that l_1 sparsified models of generic turbulent flows will have good long-term stability (Schlegel & Noack 2015), but the present results constitute evidence that this is realistically possible on a non-trivial problem. Finally, we have observed in animations of the reconstructed flow fields using the spatial modes and the temporal coefficients from numerical integration (see supplementary movie) that characteristic flow features, such as the erratic burst of corner vortices and the evolution of coherent structures in the shear layer bounding the main vortex, are also well reproduced by the l_1 model, providing a realistic flow reconstruction over long time horizons.

3.6. Sparsification set-up for DFT-based models

Before moving to the sparsification of DFT-based models, we briefly discuss three technicalities arising from the oscillatory nature of the temporal coefficients. As an illustrative example, we consider a small-sized model constructed with $N = 26$ modes and perform sparsification as discussed in § 2, with a relatively small regularisation weight ($\gamma = 10^{-14}$, strategy S1). The first key result is that all the entries of the constant and quadratic coefficient tensors \mathbf{C} and \mathbf{Q} are set to zero, while the linear tensor \mathbf{L} has a characteristic bidiagonal structure, shown in figure 13(a). The system identified by the




FIGURE 13. (a) Entries of the linear tensor \mathbf{L} identified by the unconstrained regression. (b) Singular values σ_i of the full database matrix $\Theta(\mathbf{A})$ of (2.10) (red circles), and of the reduced matrix (black crosses) obtained by keeping only the subset of columns corresponding to active interactions on the three branches of figure 5. One every five singular values are shown for clarity.

regression is equivalent to a set of $N/2$ decoupled linear oscillators in the form

$$\begin{bmatrix} \dot{a}_{2l-1} \\ \dot{a}_{2l} \end{bmatrix} = \omega_l \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} a_{2l-1} \\ a_{2l} \end{bmatrix} \quad l = 1, \dots, N/2, \quad (3.10)$$

coupling pairs of temporal coefficients oscillating at the same angular frequency $\omega_l = 2\pi l/T$, with T being the observation time. The eigendecomposition of the tensor \mathbf{L} is trivial. Eigenvalues are all imaginary and come in pairs that are integer multiples of the fundamental frequency $\omega_1 = 2\pi/T$. While this result is consistent with recent ideas on Koopman operator theory (Mezić 2013), where nonlinear dynamics is modelled with a linear system of larger dimension, all information on nonlinear energetic interactions has been lost in the process since the nonlinear part of the system has been completely eliminated by the regression. This result is due to the fact that, when temporal coefficients are sine/cosine pairs, there is a column of $\Theta(\mathbf{A})$ that is exactly parallel to the target \mathbf{A}_i , since time differentiation is equivalent to a permutation of sine/cosine pairs. As pointed out in Brunton *et al.* (2019) incorporating and enforcing known flow physics is a challenge and opportunity for machine learning algorithms. In order to address this first aspect, we introduce a physically motivated approach based on considerations of the time averaged energy budget of system (2.7). Since the temporal coefficients have zero mean and are uncorrelated in time, we should obtain

$$L_{ii}\overline{a_i a_i} + \sum_{j=1}^N \sum_{k=1}^N Q_{ijk}\overline{a_i a_j a_k} = 0, \quad i = 1, \dots, N \quad (3.11)$$

i.e. only the diagonal element of the linear term participates in the mean power budget. Hence, for the sparsification of DFT-based models we use a modified database matrix that only contains the column associated with the diagonal part of the linear term.

The second aspect is that for DFT models the database matrix $\Theta(\mathbf{A})$ is not full rank and some of the columns of this matrix are linearly dependent. In this case, the LASSO is known to select one column at random (according to the particular ordering of the columns) and sets to zero regression coefficients of the other linearly dependent columns (Tibshirani 2013; Hastie, Tibshirani & Wainwright 2015). Machine learning techniques often come without guarantees for robustness (Brunton *et al.* 2019), implying that physical insight obtained with these tools might be questionable. To avoid this problem, we

constructed a reduced database matrix $\Theta(\mathbf{A})$ containing only columns corresponding to the interactions on the three branches of figure 5. The reduced database matrix is full rank, as can be seen in figure 13(b), showing the singular values of the full database matrix defined by (2.10) and of the reduced matrix. The important consequence is that the solution of the LASSO problem (2.12) is unique (Tibshirani 1996), and can be thus compared with the available physical knowledge of scale interactions in turbulent flows. This reduction is not strictly necessary, since the l_1 regression identifies this pattern anyway for fairly small regularisation weights. However, this has the advantage that the computational complexity of sparsifying the entire Galerkin model only grows as $O(N^2)$ instead of $O(N^3)$, as for POD models, because the reduced database matrix contains a number of interactions equal to $q = 2(N + 1)$ at most. More importantly, because of the greatly reduced number of free coefficients cross-validation techniques to avoid over-fitting become unnecessary.

The third aspect of DFT-based models is that, as anticipated, the number of modes is not uniquely defined by the energy resolution but depends on the overall observation time. Long observation times would be beneficial to reach statistical significance but would result in low-energy/low-frequency modes that do not contribute significantly to the overall dynamics. In practice, we have divided the original dataset into M partitions and performed DFT for each of them separately. Then, we stacked vertically the modal acceleration matrices and the reduced database matrices from the partitions and solved (2.12) for a common coefficients vector.

3.7. Sparsification of DFT-based models

We now move to the sparsification of DFT-based models. We focus primarily on the structure of energy interactions identified by the regression and leave long-term temporal stability considerations aside. In fact, the DFT produces modal structures by assuming *a priori* their temporal behaviour, i.e. harmonic motion, and the meaning of a time-domain analysis is thus conceptually unclear.

We introduce the modified density ρ_{DFT} spanning the range $[0, 1]$ and representing the number of active coefficients with respect to the total number of active interactions on the three branches of figure 5. For large models, the approximation $\rho_{DFT} \approx 2/3\rho N$ can be used. In figure 14(a), sparsification curves for three models obtained with observation times $T = 10, 30$ and 50 (with $M = 15, 5$ and 3 partitions of the full dataset, respectively), at full energy resolution, are reported. Strategy S1, where the regularisation weight is maintained constant for all modes is used. We observe that the error decreases monotonically with the observation time. This is a consequence of the larger number of frequencies that interact quadratically to reconstruct the original DNS acceleration data. For the larger model obtained at $T = 50$, 70% of the triadic interactions can be pruned with no major effects on the overall prediction error. If the full coefficient tensor \mathbf{Q} is considered, this correspond to a remarkably low density of 0.0015. Figure 14(b) shows the sparsification curves for models obtained with observation time $T = 30$, for three different energy resolutions, $e(n) = 0.9, 0.95$ and 0.99 . Interestingly, we notice that the curves do not present a plateau for high densities as opposed to the full resolution mode shown in panel (a) and the POD sparsification curves of figure 8. This is the combined effect of the dramatic decrease in the number of modes at lower energy resolutions (see table 1) and the inherent efficient description of energy interactions in DFT-based models compared to POD.

We now compare strategies S1 and S2 on the full resolution model obtained with observation time $T = 30$. Results of this analysis are reported in figure 15. Panels (a–c) and (d–f) are obtained with the strategy S1 and S2, respectively. Panel (a) shows the tensor $\hat{\mathbf{y}}$,




FIGURE 14. (a) Sparsification curves for models obtained by three different observation times and resolving 100% of the kinetic energy. (b) Sparsification performed with $T = 30$ with three different energy resolutions $e(n)$.

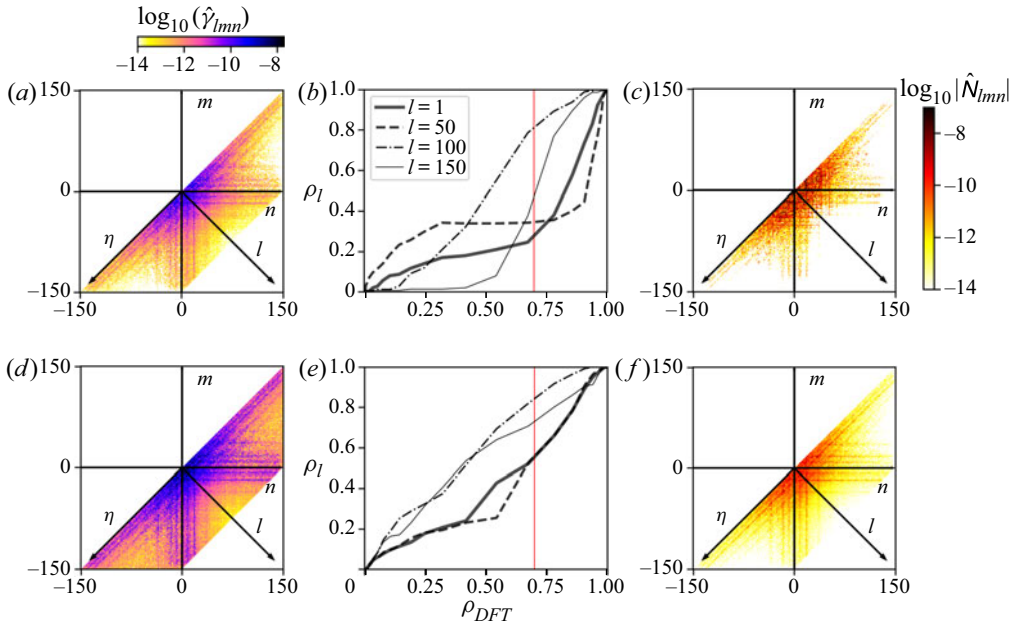


FIGURE 15. (a)–(c) Strategy S1; (d)–(f) strategy S2. Panels (a,d) show the distribution of $\hat{\gamma}$. Panels (b,e) show the trend of the modal density ρ_l against the global density ρ_{DFT} for four different modes in different parts of the spectrum. Panels (c,f) show the energy interaction tensor $\hat{\mathbf{N}}^s$ of the sparsified system.

obtained by processing and visualising the full tensor $\boldsymbol{\gamma}$ using the same technique utilised for the interaction tensor $\hat{\mathbf{N}}$ in figure 5. Panel (b) shows the density of individual ordinary differential equations for a selected number of modal structures as a function of the overall model density ρ_{DFT} , while the sparsified interaction tensor $\hat{\mathbf{N}}^s$ for $\rho_{DFT} = 0.7$ is shown in panel (c). When the regularisation weight is maintained constant, the sparsification pattern emerging from the tensor $\hat{\boldsymbol{\gamma}}$ follows the distribution of the mean energy transfer rate of figure 7(a). In particular, despite the signature of non-locality is still visible in the pattern, the sparsification is highly skewed across the spectrum because the equations for

high-frequency modes are excessively sparsified for moderate penalisation as opposed to those of low-frequency, high-energy modes.

This behaviour is better seen in the individual density curves in panel (b). Specifically, the density ρ_l of the last mode pair ($l = 150$) drops quite pronouncedly to much lower density than average at $\rho_{DFT} \approx 0.5$. Panel (d) shows the sparsification pattern obtained with the second strategy. We observe that, in this case, the interactions are retained according to their relative strength producing a sparsification pattern that follows the relative energy transfer rate reported in figure 7(b). This results in a more balanced sparsification across the spectrum, where the modal density ρ_l decreases more uniformly for all modes as the global density is decreased, as shown in panel (e). The mean energy transfer rate of the models sparsified using the two strategies, with $\rho_{DFT} = 0.7$, is reported in panels (c, f). Globally, the structure and intensity of energy interactions is preserved by the LASSO, although strategy S1 has more aggressively sparsified the high-index modes and truncated the equations of the last five pairs of modes. Although not shown here, all DFT models, regardless the strategy, have similar energy conservation properties as for POD modes, as illustrated in figure 11.

As a final remark, we have observed in sparsification of larger DFT models that, although the LASSO is able to successfully identify the dominant subset of energy interactions, the complexity of the optimisation problem makes an accurate reconstruction of the numerical values of the system coefficients challenging. This is due to the spectral properties of the database matrix Θ which deteriorate as the number of modes considered grows (Cordier *et al.* 2010). A potential solution to this issue would be to use elastic-net regression (Friedman *et al.* 2008) which combines an l_1 term with an l_2 (Tikhonov) penalisation. This would provide a better trade-off between sparsification and stability of the reconstructed coefficients.

4. Concluding remarks

In this paper, we applied recent data-driven methods for the identification of sparse dynamical systems to sparsify nonlinear triadic interactions in projection-based reduced-order models of two-dimensional unsteady flows. Our work is motivated by established knowledge of scale interactions in turbulence, whereby the dynamics at a certain length scale depends most prominently on a subset of other length scales. Computationally, our methodology is based on l_1 -based regression methods and is scalable to large models defined by hundreds of modal structures. These methods are used to recast the problem of identifying relevant triadic interactions into a convex optimisation problem for which scalable, efficient solvers can be used. The overarching aim is to develop large reduced-order models covering a wide range of length scales, but where computational efficiency and physical interpretability have been preserved by pruning weak triadic interactions.

In this analysis we considered two-dimensional lid-driven cavity flow at Reynolds number $Re = 2 \times 10^4$. We generated two families of reduced-order models by Galerkin projection of the Navier–Stokes equations onto the subspace spanned by proper orthogonal decomposition and discrete Fourier transform modes. The goal was to understand the role of the subspace utilised for projection on the structure and sparsity of energy interactions between modes. As discussed in Brunton *et al.* (2019) an open problem in applying machine learning algorithms to fluids problems is to successfully incorporate known flow physics. In our case, we have observed that for DFT-based models, it has become necessary to manually modify the database matrix in order to ensure the uniqueness of the solution

and preserve the full nonlinear character of modal interactions. The analysis of the average energy transfer rates between modal structures has shown that, for both the POD- and DFT-based models, a subset of most relevant interactions exists, in agreement with the established picture of scale interactions in two-dimensional flows. This is an *a posteriori* feature of solutions of the equations and not an *a priori* property of the evolution equations. In fact, the model coefficients identified by the Galerkin projection do not have a particular structure and are typically different from zero. Our results show that, in both cases, there exists a sweet spot on the ρ - ϵ curve where the sparsification approach recovers correctly this subset, with little effect on the prediction accuracy. The models also preserve to a good degree of accuracy the non-local nature of triadic interactions and the conservation properties of the convective term of the Navier–Stokes equations. In principle, energy conservation could be enforced exactly (Loiseau & Brunton 2018), although we have not found this to be necessary to obtain satisfactory temporal stability characteristics. In fact, unlike dense models obtained directly from the projection, the l_1 sparsified POD-based models have excellent long-term stability properties. Numerical integration shows that trajectories generated by these models remain in the same area of state space occupied by the DNS projections. The average integral fluctuation kinetic energy and the distribution of energy across modes are also reproduced fairly well. Sparse models constructed by a naive procedure where weak interactions are pruned, referred to as greedy approach in the paper, do not enjoy the same robustness and have worse performance than the dense models. This indicates that, once coefficients corresponding to weak interactions have been shrunk to zero by the l_1 penalisation, re-balancing the remaining coefficients with the least-squares term in the LASSO problem (2.12) is key to preserve accuracy.

We have also observed that the effectiveness of the sparsification grows with the number of modes utilised in the projection (energy resolution). This is a result of the non-local nature of scale interactions in two-dimensional flows, where the dynamics of small-scale features is dominated by the advection of the large modes, rather than by the small-scale/small-scale nonlinearity. The interesting consequence is that, while the total number of quadratic interactions grows cubically with the number of modes, the number of relevant interactions does not grow as quickly. Hence, our expectation is that sparsification becomes more effective as the Reynolds number increases, as a result of the increased scale separation. This would also reduce, in relative terms, the computational costs required for propagating the model forward in time. In this paper, we have not, admittedly, explored in enough detail the role of sparsification on the reduction of computational costs, as these are highly influenced by implementation details and code optimisations. For instance, the sparsified quadratic tensor \mathbf{Q}^* could be efficiently stored and evaluated using sparse matrix techniques. Characterising physical and computational properties of l_1 models as a function of the Reynolds number in two- and three-dimensional turbulent flows is therefore an interesting avenue for future work.

A major difference between the two decompositions considered in this work is that energy interactions between triads of DFT modes are highly localised in modal space, as a result of the oscillatory nature of the temporal coefficients. On the other hand, for energy-optimal POD modes, temporal coefficients contain a wider range of frequencies and energy transfers are thus inevitably scattered in modal space. Consequently, the number of active interactions only grows as $O(N^2)$ for DFT models, rather than as $O(N^3)$ for POD models. Therefore, we conclude that the sparsity of energy interactions is not necessarily invariant when analysed on different subspaces. In practice, the favourable $O(N^2)$ scaling observed for DFT modes is appealing for the construction of large yet interpretable models, covering a wide range of spatio-temporal scales. However, the meaning and practical utility of temporal integration of DFT-based models (but also

models constructed from other decompositions producing purely oscillatory modes), where the solution structure has been assumed *a priori*, remains conceptually unclear. In summary, the stark difference between DFT- and POD-based models suggests that it might be possible to develop a modal decomposition that identifies a set of maximally independent structures, e.g. where the resulting quadratic coefficient tensor \mathbf{Q} or the average interaction tensor \mathbf{N} are maximally sparse. However, this appears to be a nonlinear optimisation problem, with the associated convergence and uniqueness issues. Work in this direction has been recently initiated by Schmidt (2020).

Acknowledgements

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.

Supplementary movies

A supplementary movie is available at <https://doi.org/10.1017/jfm.2020.707>.

REFERENCES

- 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.
- BLUM, A. L. & LANGLEY, P. 1997 Selection of relevant features and examples in machine learning. *Artif. Intell.* **97** (1–2), 245–271.
- BOPANA, V. B. L. & GAJJAR, J. S. B. 2010 Global flow instability in a lid-driven cavity. *Intl J. Numer. Meth. Fluids* **62** (8), 827–853.
- BOTELLA, O. & PEYRET, R. 1998 Benchmark spectral results on the lid-driven cavity flow. *Comput. Fluids* **27** (4), 421–433.
- 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., NOACK, B. R. & KOUMOUTSAKOS, P. 2019 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.
- 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. USA* **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.
- CHEN, K. K., TU, J. H. & ROWLEY, C. W. 2012 Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses. *J. Nonlinear Sci.* **22** (6), 887–915.
- 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., BASDEVANT, C. & SAGAUT, P. 2005 Calibrated reduced-order POD-Galerkin system for fluid flow modelling. *J. Comput. Phys.* **207** (1), 192–220.

- 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.
- 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.
- FLETCHER, C. A. J. 1984 *Computational Galerkin Methods*, 1st edn. Springer.
- FRIEDMAN, J., HASTIE, T. & TIBSHIRANI, R. J. 2008 *The Elements of Statistical Learning*. Springer.
- HASTIE, T., TIBSHIRANI, R. & WAINWRIGHT, M. 2015 *Statistical Learning with Sparsity: The LASSO and Generalizations*. Chapman and Hall/CRC.
- JIMÉNEZ, J. 2018 Coherent structures in wall-bounded turbulence. *J. Fluid Mech.* **842**, P1.
- JOVANOVIĆ, M. R., SCHMID, P. J. & NICHOLS, J. W. 2014 Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **26** (2), 024103.
- KAISER, E., KUTZ, J. N. & BRUNTON, S. L. 2018 Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proc. Math. Phys. Engng Sci.* **474** (2219), 20180335.
- 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.
- LOISEAU, J. C. & BRUNTON, S. L. 2018 Constrained sparse Galerkin regression. *J. Fluid Mech.* **838**, 42–67.
- LUMLEY, J. L. 1970 *Stochastic Tools in Turbulence*. Academic Press.
- LUMLEY, J. L. 1979 Computational modeling of turbulent flows. In *Advances in Applied Mechanics*, vol. 18, pp. 123–176. Elsevier.
- MENDEZ, M. A., BALABANE, M. & BUCHLIN, J.-M. 2019 Multi-scale proper orthogonal decomposition of complex fluid flows. *J. Fluid Mech.* **870**, 988–1036.
- MEZIĆ, I. 2013 Analysis of fluid flows via spectral properties of the Koopman operator. *Annu. Rev. Fluid Mech.* **45**, 357–378.
- NAIR, A. G., BRUNTON, S. L. & TAIRA, K. 2018 Networked-oscillator-based modeling and control of unsteady wake flows. *Phys. Rev. E* **97**, 063107.
- NAIR, A. G. & TAIRA, K. 2015 Network-theoretic approach to sparsified discrete vortex dynamics. *J. Fluid Mech.* **768**, 549–571.
- NEWMAN, M. 2018 *Networks*. Oxford University Press.
- NOACK, B. R. & ECKELMANN, H. 1994 A global stability analysis of the steady and periodic cylinder wake. *J. Fluid Mech.* **270**, 297–330.
- 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-Equil. 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.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., *et al.* 2011 Scikit-learn: machine learning in Python. *J. Machine Learning Res.* **12**, 2826–2830.
- 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.
- PERRET, L., COLLIN, E. & DELVILLE, J. 2006 Polynomial identification of POD based low-order dynamical system. *J. Turbul.* **7**, N17.
- POLIASHENKO, M. & AIDUN, C. K. 1995 A direct method for computation of simple bifurcations. *J. Comput. Phys.* **121** (2), 246–260.

- POPE, S. B. 2001 *Turbulent Flows*. Cambridge University Press.
- RAMIREZ, C., KREINOVICH, V. & ARGAEZ, M. 2013 Why l_1 is a good approximation to l_0 : a geometric explanation. *J. Uncertain Syst.* **7** (3), 203–207.
- 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.
- ROWLEY, C. W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
- 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. [arXiv:2002.04146](https://arxiv.org/abs/2002.04146).
- SCHMIDT, M. & LIPSON, H. 2009 Distilling free-form natural laws from experimental data. *Science* **324** (5923), 81–85.
- SIEBER, M., PASCHEREIT, C. O. & OBERLEITHNER, K. 2016 Spectral proper orthogonal decomposition. *J. Fluid Mech.* **792**, 798–828.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structure. Part I, II, III. *Q. Appl. Maths* **3**, 583.
- 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**, 4013–4041.
- TAIRA, K., NAIR, A. G. & BRUNTON, S. L. 2016 Network structure of two-dimensional decaying isotropic turbulence. *J. Fluid Mech.* **795**, R2.
- 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.
- THOMAS, V. L., FARRELL, B. F., IOANNOU, P. J. & GAYME, D. F. 2015 A minimal model of self-sustaining turbulence. *Phys. Fluids* **27** (10), 105104.
- THOMAS, V. L., LIEU, B. K., JOVANOVIĆ, M. R., FARRELL, B. F., IOANNOU, P. J. & GAYME, D. F. 2014 Self-sustaining turbulence in a restricted nonlinear model of plane Couette flow. *Phys. Fluids* **26** (10), 105112.
- 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.
- XIE, X., MOHEBUJAMAN, M., REBHOLZ, L. G. & ILIESCU, T. 2018 Data-driven filtered reduced order modeling of fluid flows. *SIAM J. Sci. Comput.* **40** (3), B834–B857.
- YEUNG, P. K., BRASSEUR, J. G. & WANG, Q. 1995 Dynamics of direct large-small scale couplings in coherently forced turbulence: concurrent physical- and fourier-space views. *J. Fluid Mech.* **283**, 43–95.
- ZHANG, L. & SCHAEFFER, H. 2019 On the convergence of the SINDy algorithm. *Multiscale Model. Simul.* **17** (3), 948–972.