Dynamic mode decomposition of numerical and experimental data

PETER J. SCHMID†

Laboratoire d'Hydrodynamique (LadHyX), CNRS-École Polytechnique, 91128 Palaiseau, France

The description of coherent features of fluid flow is essential to our understanding of fluid-dynamical and transport processes. A method is introduced that is able to extract dynamic information from flow fields that are either generated by a (direct) numerical simulation or visualized/measured in a physical experiment. The extracted dynamic modes, which can be interpreted as a generalization of global stability modes, can be used to describe the underlying physical mechanisms captured in the data sequence or to project large-scale problems onto a dynamical system of significantly fewer degrees of freedom. The concentration on subdomains of the flow field where relevant dynamics is expected allows the dissection of a complex flow into regions of localized instability phenomena and further illustrates the flexibility of the method, as does the description of the dynamics within a spatial framework. Demonstrations of the method are presented consisting of a plane channel flow, flow over a two-dimensional cavity, wake flow behind a flexible membrane and a jet passing between two cylinders.

1. Introduction and motivation

The accurate description of the disturbance behaviour in complex geometries poses a great challenge to numerical simulations and physical experiments, as well as to the computational algorithms that extract and quantify this behaviour. At the same time, many industrial applications, such as flow in a combustion chamber, could greatly benefit from a more thorough understanding of the underlying transition and instability mechanisms.

Global stability analyses for flows in complex geometries are becoming more commonplace, but the resulting large stability matrix sizes have put considerable strain on computational resources. Direct methods, the method of choice for simple problems, become prohibitively expensive, and iterative schemes have thus to be employed to extract the global stability modes. The Arnoldi scheme (see Edwards *et al.* 1994) has been particularly successful in this respect. It is based on an approximation of the high-dimensional system matrix by projecting it onto a lower-dimensional Krylov subspace. In this way, the dominant eigenvalues (and corresponding eigenvectors) of the full system can be computed rather efficiently. Convergence can be improved in various ways by shifting, mapping, restarting, locking and purging techniques (Lehoucq & Scott 1997), but in most cases this is accomplished at the expense of computational efficiency due to the necessity of additional matrix inversions.

† Email address for correspondence: peter@ladhyx.polytechnique.fr

The extraction of dynamical features by a global stability analysis has remained a tool that is nearly exclusively applied to numerical simulations. This is because the respective algorithms require the system matrix of the underlying flow in order to build a sequence of (artificial) flow fields upon which the convergence of the numerical method relies. In physical experiments this system matrix is not available, and whereas a subroutine call of the right-hand side is straightforward to accommodate in a numerical simulation, the same is not true for experiments. Rather, in experimental set-ups, the only input that is readily available are the flow fields themselves, either in form of particle image velocimetry (PIV) measurements or in form of visualizations of a passive tracer. As a consequence, if coherent structures are to be identified from experimental data, algorithms need to be designed that rely on these measurements only.

A common technique for identifying coherent structures is the proper orthogonal decomposition (POD) method (Lumley 1970; Sirovich 1987; Berkooz, Holmes & Lumley 1993). This method is capable of extracting information from snapshots of the flow field and is thus applicable to experimental data (see Herzog 1986; Bonnet *et al.* 1994; Noack *et al.* 2003). The method determines the most energetic structures by diagonalizing the spatial correlation matrix computed from the snapshots. Two major drawbacks that are tacitly acknowledged by employing this method are associated with this technique: (i) the energy may not in all circumstances be the correct measure to rank the flow structures, and (ii) due to the choice of second-order statistics as a basis for the decomposition, valuable phase information is lost. The first shortcoming has been widely recognized, and an explanation for the existence of dynamically highly relevant but zero-energy modes has been presented by Noack *et al.* (2008). Choosing weight functions that put more emphasis on specific components of the flow field or more active regions of the flow can ameliorate the focus on the total kinetic perturbation energy. The second shortcoming is more difficult to overcome. The bi-orthogonal decomposition (BOD; see Aubry 1991; Hemon & Santi 2007) simultaneously produces temporal and spatial structures, referred to as 'chronos' and 'topos', via a singular value decomposition of the snapshot matrix; the same result can also be obtained by choosing a temporal or spatial average when computing the correlation tensor, or by using the classical POD (Lumley 1970) or snapshot POD (Sirovich 1987), respectively. The temporal structures (*chronos*) represent the eigenvectors of the spatially averaged temporal correlation matrix, while the spatial structures (*topos*) constitute the eigenvectors of the temporally averaged spatial correlation matrix. In either case, the averaging process that produces second-order statistics causes the loss of information that might be important when classifying the dynamic processes contained in the snapshots.

In this article we will propose a decomposition method that is based on snapshots of the flow only (see Schmid & Sesterhenn 2008), but that will still yield fluid structures that accurately describe the motion of the flow. In the case of a linearized flow (i.e. a flow of small perturbation about a steady base flow), the extracted modes are equivalent to the result of a global stability analysis; for a nonlinear flow, the results produce structures of a linear tangent approximation to the underlying flow and describe fluid elements that express the dominant dynamic behaviour captured in the data sequence. To delineate ourselves from a classical linear global stability analysis, we will refer to the extracted flow structures as 'dynamic modes' and to the decomposition technique as the 'dynamic mode decomposition'. This technique is at the basis of a Koopman analysis of nonlinear dynamical systems (see Lasota & Mackey 1994; Mezić 2005), which has recently been applied to large-scale simulations of a jet in crossflow by Rowley *et al.* (2009). Connections and differences to commonly

applied decompositions will be pointed out that will help put the new method into perspective with familiar techniques of describing coherent structures. We will then validate and apply this method to three different flow cases: one of them based on data generated numerically by a Navier–Stokes code, the last two based on data extracted from experimental flow measurements. We stress that in each case only the flow fields will be used in extracting relevant dynamic information.

2. Dynamic mode decomposition using flow field data

The derivation of the dynamic mode decomposition (DMD) will closely follow arguments used in iterative methods for computing solutions to linear eigenvalue or other linear algebra problems. Furthermore, readers familiar with the techniques of proper orthogonal (Lumley 1970; Sirovich 1987) or bi-orthogonal decompositions (Aubry 1991; Hemon & Santi 2007) will notice commonalities; important differences, however, do exist and will be pointed out below.

As mentioned previously, we aim at a method that is equally applicable to experimental and numerical flow field data; we thus attempt a 'matrix-free' formulation relying only on the gathered input data and ignoring any information about the underlying system matrix. In effect, we avoid a model-based approach to extracting dynamic information and, instead, concentrate on a data-based procedure.

2.1. *General description*

We will start with a general description of flow fields collected by sampling either direct numerical simulations or experimental data. Preprocessing the data may be necessary for experimental data in order to eliminate inherent measurement noise. The data shall be represented in the form of a snapshot sequence, given by a matrix $\boldsymbol{V}_1^N$,

$$\boldsymbol{V}_1^N = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \ldots, \boldsymbol{v}_N\}, \tag{2.1}$$

where $\boldsymbol{v}_i$ stands for the ith flow field. In the above definition, the subscript 1 denotes the first member of the sequence, while the superscript N denotes the last entry in the sequence, i.e. the first and last columns of the matrix $\boldsymbol{V}_1^N$, respectively. We further assume an ordered sequence of data separated by a constant sampling time Δt; the choice of the time step Δt between two consecutive snapshots will be further discussed in §2.6.

In the first step, we assume that a linear mapping $\boldsymbol{A}$ connects the flow field $\boldsymbol{v}_i$ to the subsequent flow field $\boldsymbol{v}_{i+1}$, that is,

$$\boldsymbol{v}_{i+1} = \boldsymbol{A}\boldsymbol{v}_i, \tag{2.2}$$

and that this mapping is approximately the same over the full sampling interval $[0, (N-1)\Delta t]$. If the flow fields stem from a nonlinear process, this assumption amounts to a linear tangent approximation. For slowly varying systems, a multiple-scale argument can provide a foundation for the above assumption. In the special case of a purely linear process, no approximation is invoked by assuming a constant mapping. In any case, the assumption of a constant mapping between the snapshots $\boldsymbol{v}_i$ will allow us to formulate our sequence of flow fields as a Krylov sequence (see e.g. Greenbaum 1997; Trefethen & Bau 1997).

$$\boldsymbol{V}_1^N = \{\boldsymbol{v}_1, \boldsymbol{A}\boldsymbol{v}_1, \boldsymbol{A}^2\boldsymbol{v}_1, \ldots, \boldsymbol{A}^{N-1}\boldsymbol{v}_1\}. \tag{2.3}$$

Our goal then is the extraction of the dynamic characteristics (eigenvalues, eigenvectors, pseudoeigenvalues, energy amplification, resonance behaviour, etc.) of the dynamical process described by $\boldsymbol{A}$ based on the sequence $\boldsymbol{V}_1^N$.

As the number of snapshots increases and the data sequence given by $\boldsymbol{V}_1^N$ captures the dominant features of the underlying physical process, it is reasonable to assume that, beyond a critical number of snapshots, the vectors given by (2.2) become linearly dependent. In other words, adding further flow fields $\boldsymbol{v}_i$ to the data sequence will not improve the vector space spanned by $\boldsymbol{V}_1^N$. When this limit is reached, we can express the vector $\boldsymbol{v}_N$ as a linear combination of the previous, and linearly independent, vectors $\boldsymbol{v}_i, i = 1, \ldots, N-1$ according to

$$\boldsymbol{v}_N = a_1 \boldsymbol{v}_1 + a_2 \boldsymbol{v}_2 + \cdots + a_{N-1} \boldsymbol{v}_{N-1} + \boldsymbol{r} \tag{2.4}$$

or in matrix form

$$\boldsymbol{v}_N = \boldsymbol{V}_1^{N-1} \boldsymbol{a} + \boldsymbol{r} \tag{2.5}$$

with $\boldsymbol{a}^{\mathrm{T}} = \{a_1, a_2, \ldots, a_{N-1}\}$ and $\boldsymbol{r}$ as the residual vector. We continue by following Ruhe (1984) and write

$$\boldsymbol{A}\{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \ldots, \boldsymbol{v}_{N-1}\} = \{\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4, \ldots, \boldsymbol{v}_N\} = \{\boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4, \ldots, \boldsymbol{V}_1^{N-1} \boldsymbol{a}\} + \boldsymbol{r}\boldsymbol{e}_{N-1}^{\mathrm{T}} \tag{2.6}$$

or in matrix form

$$\boldsymbol{A}\boldsymbol{V}_1^{N-1} = \boldsymbol{V}_2^N = \boldsymbol{V}_1^{N-1} \boldsymbol{S} + \boldsymbol{r}\boldsymbol{e}_{N-1}^{\mathrm{T}} \tag{2.7}$$

with $\boldsymbol{e}_{N-1} \in \mathbb{R}^{N-1}$ as the $(N-1)$th unit vector.

A simple calculation shows that the matrix $\boldsymbol{S}$ is of companion type with

$$\boldsymbol{S} = \begin{pmatrix} 0 & & & & a_1 \\ 1 & 0 & & & a_2 \\ & \ddots & \ddots & & \vdots \\ & & 1 & 0 & a_{N-2} \\ & & & 1 & a_{N-1} \end{pmatrix}, \tag{2.8}$$

whose subdiagonal entries reflect the fact that, by design, the vector in the ith column of $\boldsymbol{V}_2^N$ is identical to the vector in the $(i+1)$th column of $\boldsymbol{V}_1^{N-1}$ for $i = 1, \ldots, N-2$. The only unknowns in $\boldsymbol{S}$ are the coefficients $\{a_1, a_2, \ldots, a_{N-1}\}$ which constitute the above-mentioned $(N-1)$-component linear representation of the last sample $\boldsymbol{v}_N$ in terms of the previous samples $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \ldots, \boldsymbol{v}_{N-1}\}$.

The eigenvalues of $\boldsymbol{S}$ then approximate some of the eigenvalues of $\boldsymbol{A}$. The well-known Arnoldi method (see e.g. Greenbaum 1997; Trefethen & Bau 1997) is closely related to the decomposition above but successively orthogonalizes the vectors of $\boldsymbol{V}_1^N$ resulting in a decomposition of the form $\boldsymbol{A}\boldsymbol{Q} \approx \boldsymbol{Q}\boldsymbol{H}$ with $\boldsymbol{V}_1^{N-1} = \boldsymbol{Q}\boldsymbol{R}$ and $\boldsymbol{H} = \boldsymbol{R}\boldsymbol{S}\boldsymbol{R}^{-1}$ as a Hessenberg matrix. Again, the eigenvalues of $\boldsymbol{H}$ approximate some of the eigenvalues of $\boldsymbol{A}$. In practice, the reduction of $\boldsymbol{A}$ to Hessenberg form by the Arnoldi method is not accomplished by a simple QR-decomposition of $\boldsymbol{V}_1^{N-1}$, but rather by a sequence of projections onto successive Krylov subspaces. This yields a more stable algorithm, but for these projections the matrix $\boldsymbol{A}$ has to be available which makes the classical Arnoldi method unattractive for our purposes. Rather, we shall contend with less favourable stability (and convergence) properties of our algorithm in order to gain a numerical technique that is exclusively based on flow fields and is thus equally applicable to experimental data and large-scale numerical simulations.

The computation of $\boldsymbol{S}$ then proceeds as follows: the last element of a given data sequence $\boldsymbol{v}_N$ is expressed as a linear combination of the previous elements of the sequence as stated in (2.5) whose least-squares solution, for a full-rank matrix $\boldsymbol{V}_1^{N-1}$,

is given by

$$a = R^{-1} Q^H v_N, \tag{2.9}$$

with $QR = V_1^{N-1}$ as the economy-size QR-decomposition of the data sequence V_1^{N-1}. The $(N-1)$-component vector a then forms the last column of the companion matrix S.

Even though the above decomposition based on a companion matrix S is mathematically correct and is often used to prove convergence properties of the full Arnoldi method (Ruhe 1984), a practical implementation yields an ill-conditioned algorithm that is often not capable of extracting more than the first or first two dominant dynamic modes. This is particularly true when the data stem from an experiment and are contaminated with noise and other uncertainties. For this reason, we choose a more robust implementation that results in a 'full' matrix $\tilde{S}$ – related to S via a similarity transformation. Robustness is achieved by a preprocessing step using a singular value decomposition of the data sequence $V_1^{N-1} = U \Sigma W^H$. Substituting the singular value decomposition $U \Sigma W^H$ into (2.7) and rearranging the resulting expression we obtain $U^H A U = U^H V_2^N W \Sigma^{-1} \equiv \tilde{S}$. By recognizing that the matrix U contains the proper orthogonal modes of the data sequence V_1^{N-1}, the above operation amounts to a projection of the linear operator A onto a POD basis (more details about this connection will be given in §2.2). A further advantage of this operation, besides a more robust calculation of the low-dimensional representation of A, is the opportunity to account for a rank-deficiency in the data sequence V_1^{N-1} via a restriction to a limited projection basis U given by the non-zero singular values of Σ (or by singular values above a prescribed threshold). The implementation based on $\tilde{S}$ has been used throughout to obtain the results in this paper.

The modal structures are extracted from the matrix $\tilde{S}$ in a manner analogous to recovering the global modes from the eigenvectors of the Hessenberg matrix H of the standard Arnoldi method. In our case we have the following expression for the dynamic modes Φ_i:

$$\Phi_i = U y_i, \tag{2.10}$$

with y_i as the ith eigenvector of $\tilde{S}$, i.e. $\tilde{S} y_i = \mu_i y_i$, and U as the right singular vectors of the snapshot sequence V_1^{N-1}. The above decomposition method, whether in its mathematical form based on a companion matrix or in its implementation based on a full matrix, is able to extract coherent structure from a sequence of data fields only.

2.2. Relation to other decompositions

Expression (2.7) clearly shows relations to other commonly applied decomposition techniques. As mentioned previously, the Arnoldi method accomplishes the reduction of a matrix A to a smaller Hessenberg matrix by successive projections onto an orthonormalized Krylov sequence. In our case no orthogonalization step is taken; instead, a smaller system matrix is generated from the snapshots directly. This feature comes at the expense of reduced algorithmic stability and convergence properties but, in return, allows a 'model-free' application.

A method commonly applied in the analysis of atmospheric or oceanographic data is the decomposition into principal interaction patterns (PIPs) or principal oscillation patterns (POPs), a statistical method that recovers coherent patterns from the ratio of a temporal cross-correlation over one time step to the auto-correlation of the original field (Hasselmann 1988; von Storch *et al.* 1995). Starting with the hypothesis that the

data $x(t)$ are generated by a linear process of the form

$$x(t + \Delta t) = Ax(t) + \text{noise}, \tag{2.11}$$

multiplication from the right by the transpose data vector $x^H(t)$ and averaging in time yields an expression for the matrix A according to

$$A = \overline{x(t + \Delta t)x(t)^H} \; \overline{(x(t)x(t)^H)}^{-1}. \tag{2.12}$$

The first term on the right-hand side represents the temporal cross-correlation over a characteristic time interval Δt, while the second term denotes the inverse of the auto-correlation of the data field. The eigenvectors of A are referred to as POPs. A variant of this techniques is known as linear inverse modelling (LIM; see e.g. Penland & Magoriam 1993; DelSole & Hou 1999). It is based on a common system identification approach borrowed from signal processing where the ratio of the cross-correlation of two signals and the autocorrelation of one of the signals produces a mapping P between the two signals. In our notation, we have

$$S = \left[(V_1^{N-1})^H V_1^{N-1} \right]^{-1} (V_1^{N-1})^H V_2^N. \tag{2.13}$$

Even though related to POPs, the above expression does not attempt to recover A from $AV_1^{N-1} = V_2^N$ by using correlation techniques but, rather, it tries to extract S from $V_2^N = V_1^{N-1}S$ by a left multiplication with the transpose data matrix $(V_1^{N-1})^H$. In this sense, LIM can be thought of as a POP analysis, performed in the space of the data snapshots. In addition, the solution of S is not accomplished by a QR-decomposition, but is based on the formulation of a normal equation. The expression in (2.13) $[(V_1^{N-1})^H V_1^{N-1}]^{-1}(V_1^{N-1})^H$ is recognized as the pseudo-inverse of the data matrix V_1^{N-1}. The resulting matrix S is of companion type for a full-rank matrix V_1^{N-1}.

As mentioned earlier, a popular method for the extraction of coherent fluid structures from a sequence of snapshots is the POD (e.g. Lumley 1970; Sirovich 1987), which is also known under the names empirical orthogonal eigenfunction (EOF) analysis, principal component analysis (PCA) or Karhunen–Loève (KL) decomposition. It is based on an energy ranking of orthogonal structures computed from a correlation matrix of the snapshots. In terms of the notation above, the POD modes are contained in the singular value decomposition of the matrix $V_1^{N-1} = U\Sigma W^H$, where the matrix U contains the spatial structures (referred to as topos) and the matrix W contains the temporal structures (referred to as chronos) (see Aubry 1991; Hemon & Santi 2007). The energy ranking is given by the singular values contained in Σ. A simple analysis, given previously, shows the relation between the dynamic mode decomposition (2.7) and the POD modes (in this case the topos modes). We have

$$\tilde{S} = U^H A U, \tag{2.14}$$

which demonstrates that the low-dimensional system matrix $\tilde{S}$ is obtained by correlating the POD modes U with the structures shifted over one time step Δt, expressed as AU. This relation also demonstrates that the matrix $\tilde{S}$ contains more information about the temporal evolution of the underlying process than the time-averaged POD modes given by U. The computation of the matrix $\tilde{S}$ by the above expression (2.14) is only feasible for numerical simulations where arbitrary initial conditions can be evolved over one time-step using the system matrix A; for experiments, the above expression cannot be put into practice. Instead, the matrix $\tilde{S}$ has to be computed from the data sequence directly as outlined above.

2.3. *Application to subdomains and low-dimensional representations*

A great advantage of the dynamic mode decomposition is its detachment from an underlying model in the form of a system matrix $\boldsymbol{A}$. This feature also allows the processing of subdomains, i.e. the forming of a vector sequence $\boldsymbol{V}_1^N$ which contains only data from a smaller part of the full computational or experimental domain. In this manner one can focus on smaller regions where dynamically interesting phenomena are expected or observed. This advantage is also shared by the POD method; a standard global mode analysis, on the other hand, has difficulties with processing subdomains because boundary conditions have to be imposed on the subdomains (e.g. Åkervik *et al.* 2007) which, when carelessly chosen, can contaminate the computed global modes.

A rather extreme example of subdomains which is common in experiments are lower-dimensional slices of a higher-dimensional flow field. For example, in experiments it is commonplace to visualize and measure three-dimensional flow fields by extracting two-dimensional slices. If the projection from three-dimensional to two-dimensional flow fields is denoted by $\mathscr{P}_{3D \to 2D}$ we have the following sequence of snapshots

$$\mathscr{P}_{3D \to 2D} \left\{ \boldsymbol{v}_0, \boldsymbol{A}_{3D} \boldsymbol{v}_0, \boldsymbol{A}_{3D}^2 \boldsymbol{v}_0, \ldots \right\}, \tag{2.15}$$

where the operator $\boldsymbol{A}_{3D}$ denotes the three-dimensional propagator from one three-dimensional data field to the next over a time interval Δt. Proceeding with the dynamic mode decomposition, we observe that the projector $\mathscr{P}_{3D \to 2D}$ is only influencing the spatial structure of the flow fields and the spatial structure of the dynamic modes, but not their temporal dynamics. The temporal dynamics (frequencies) of the three-dimensional fluid flow can thus be captured by processing the two-dimensional data slices, provided that the two-dimensional data fields contain a non-zero projection of the dynamically dominant three-dimensional structure. The associated dynamic modes, which will necessarily be two-dimensional in accordance with the two-dimensional input data, approximate the projection of a three-dimensional dynamic structure onto the slicing plane. Other examples of subdomain decomposition are signals extracted by a rake of hot-wire probes or the extraction of near- or far-field processes in open flows; in these cases the treatment of subdomains may bring significant advantages over a decomposition of the entire flow field.

2.4. *Spatial stability analysis*

In hydrodynamic stability theory we distinguish between a temporal and a spatial analysis (see e.g. Schmid & Henningson 2001). Whereas the temporal analysis treats the evolution in time of spatially wave-like solutions, the spatial approach is concerned with the growth or decay of time-harmonic (or generally time-dependent) perturbations up- and downstream of their location of origin. The spatial growth of disturbances shedding from a localized roughness element would be a typical flow configuration where a spatial stability analysis would be appropriate. The distinction between a temporal and spatial framework conceptually carries over to a global stability analysis, even though it has yet to be explored more fully for complex two- and three-dimensional flows.

Because a system matrix $\boldsymbol{A}$, the mapping from one snapshot to the next, is neither assumed nor formed for the dynamic mode decomposition, a spatial analysis of the sampled flow fields should be equally feasible. In this case we reorganize the flow fields such that they present a snapshot sequence in space. In effect, this procedure
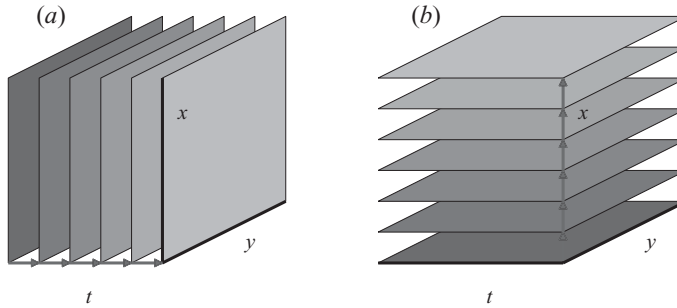
FIGURE 1. Sketch of the snapshot sequence for a temporal (*a*) and spatial (*b*) stability analyses.

amounts to slicing differently through a spatio-temporal data set (see figure 1). The low-dimensional full-rank matrix $\tilde{\mathbf{S}}$, extracted from a spatially ordered data sequence in the same above-described manner, then represents the mapping from a flow field at one spatial location to a flow field at the next spatial location. The eigenvalues of $\tilde{\mathbf{S}}$ consequently give information about the spatial dynamics of the underlying flow.

2.5. *Convergence behaviour and validation*

For the case of a linear process, as the number of flow fields in our sequence is increased, we expect convergence of the eigenvalue values of $\tilde{\mathbf{S}}$ towards some of the eigenvalues of $\mathbf{A}$. As is typical for iterative eigenvalue algorithms, the eigenvalues at the extreme ends of the spectrum of $\mathbf{A}$ are isolated first. The overall convergence can be monitored by evaluating the residual defined in (2.7) as the iterations proceed and plotting its L_2-norm. A demonstration of the convergence of the dynamic mode decomposition based on the residual measure is given in the next section.

Another interesting and more graphical convergence measure for iterative techniques are the convergence lemniscates (Trefethen & Bau 1997). The lemniscates are defined as the contour levels of the matrix polynomial $\mathscr{P}(\mathbf{A})$ representing the true system dynamics given by matrix $\mathbf{A}$ and represented in the data sequence (2.3) (see Trefethen & Bau 1997, for more details). In our case, the polynomial coefficients are given by the coefficient vector $\mathbf{a}$ using (2.9); taking advantage of the companion shape of $\mathbf{S}$, the zeros of this polynomial coincide with the eigenvalues of $\mathbf{S}$ or $\tilde{\mathbf{S}}$. As the number of snapshots, and thus the degree of the approximating polynomial $\mathscr{P}$, increases, the contours of $\mathscr{P}(z)$ in the complex z-plane will tightly encircle the eigenvalues of $\mathbf{A}$. These contours therefore present a graphical way of monitoring the convergence towards the global spectrum for a linear process. We can also observe the above-mentioned convergence characteristic towards extreme eigenvalues. To demonstrate the convergence behaviour of the dynamic mode decomposition, we use a numerically generated sequence of snapshots from a spectral discretization of the two-dimensional linearized Navier–Stokes equations for plane Poiseuille flow. The governing equations have been formulated as an evolution equation for the wall-normal velocity v, and 150 Chebyshev polynomials have been used to discretize the wall-normal direction; a matrix exponential has been used to advance the flow fields over a fixed time interval $\Delta t = 0.1$. After a transient period of ten time-steps, v-flow fields are sampled at intervals of $\Delta t = 0.1$. For demonstration purposes, we choose a Reynolds number of $Re = 10\,000$, based on the centreline velocity and half-channel height (Orszag 1971; Schmid & Henningson 2001). Furthermore, we take an initial perturbation that has a unit streamwise wavenumber $\alpha = 1$. This flow configuration
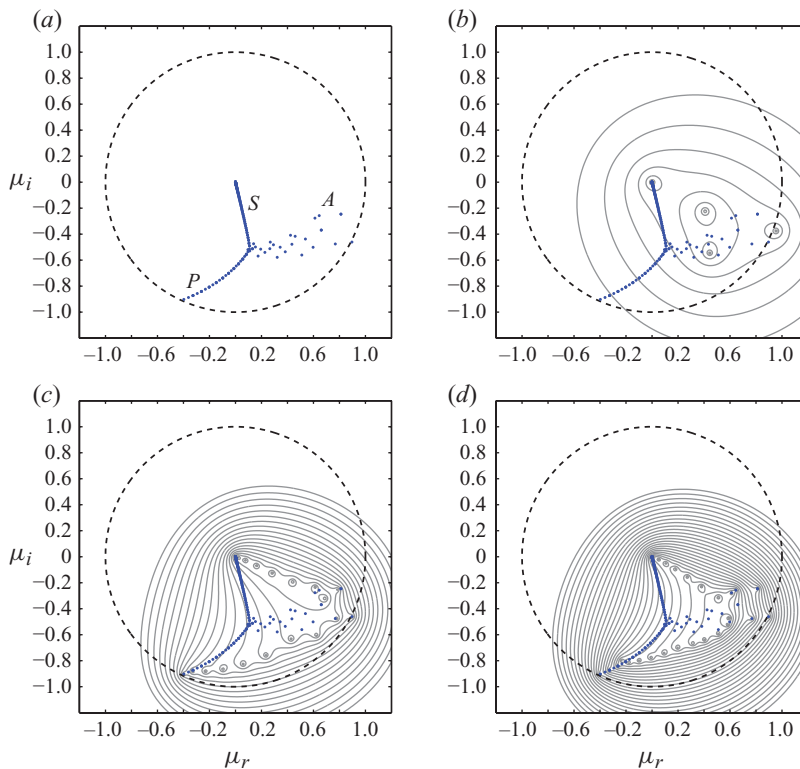
FIGURE 2. (*a*) Eigenvalues μ_i of $\tilde{\mathbf{S}}$ for plane Poiseuille flow with $Re = 10000$ and $\alpha = 1$. Because of the time-stepper format, unstable eigenvalues lie outside the unit circle. (*b*–*d*) Convergence lemniscates for a sequence of 5, 20, 40 snapshots. Contour levels range logarithmically from 10^{-30} to 1.

is often used as a benchmark problem for stability calculations and shall in our case give a first impression of the accuracy and convergence behaviour of the dynamic mode decomposition.

The results are displayed in figure 2. In figure 2(*a*), one recognizes the typical *Y*-shaped spectrum displaying the A-, P- and S-branches, but in the maybe less familiar time-stepper format. In this format, unstable eigenvalues fall outside the unit disk (marked by a black dashed circle); the closer the eigenvalues are to the origin the more damped they are. The remaining subfigures show the convergence behaviour of the dynamic mode decomposition as the number of snapshots is increased from $N = 5$ to $N = 20$ to $N = 40$. A characteristic tightening of the contours about the exact eigenvalues (marked by blue symbols) is clearly visible. The most unstable eigenvalues are identified for even a modest number of snapshots, after which the algorithm converges towards more damped modes. Saturation is reached when additional snapshots do not add information that is significantly different from the current representation. At this stage the characteristic polynomial displays a number of zeros that show no further improvement as additional snapshots are added; the algorithm then terminates.

These results can easily be contrasted to the full Arnoldi iteration by using a given number N of snapshots for the dynamic mode decomposition and by using a Krylov subspace of equal dimension N. Two cases have been considered, with $N = 50$
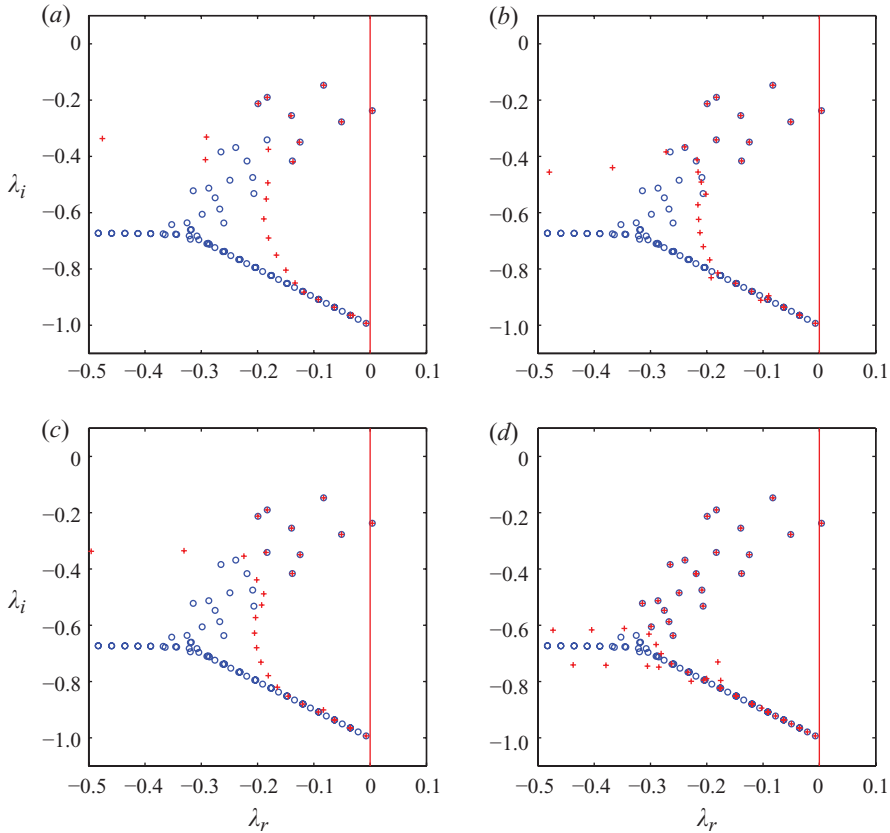
FIGURE 3. Convergence of Arnoldi (*b, d*) and dynamic mode decomposition (*a, c*) for a Krylov subspace/snapshot basis of $N = 50$ components (*a, b*) and $N = 100$ components (*c, d*) from a numerical simulation of plane Poiseuille flow at $Re = 10\,000$. The approximate eigenvalues (Ritz values) of the two methods are shown in red; the exact eigenvalues are displayed in blue.

and $N = 100$, and the results are displayed in figure 3. In this figure, the more familiar display of the stable (left) and unstable (right) complex half-plane has been chosen. This is accomplished by the logarithmic mapping of the eigenvalues according to $\lambda = \log(\mu)/\Delta t$ with Δt as the temporal separation between two consecutive snapshots. The approximate eigenvalues (also referred to as Ritz values) from the full Arnoldi method show a similar convergence pattern that is commonplace for iterative eigenvalue methods: eigenvalues at the extreme ends of the spectrum are identified by the iterative method. The number and quality of identified eigenvalues is rather similar for our case; by continuing the iterations, however, the full Arnoldi method eventually identifies the remaining eigenvalues of the problem (figure 3*d*), whereas the dynamic mode decomposition will saturate at a specific number of processed snapshots (figure 3*c*). This should not come as a surprise given the fact that the Arnoldi method can take advantage of the system matrix ***A*** to construct an orthonormal basis, whereas the dynamic mode decomposition only relies on the (nonorthogonal) snapshot sequence to accomplish the same goal. Nevertheless, the structures (and eigenvalues) identified by the dynamic mode decomposition correctly and accurately capture the temporal behaviour contained in the processed data sequence.

2.6. *Miscellaneous*

A crucial parameter in the dynamic mode decomposition is the temporal (or spatial) separation of the snapshots in the input data sequence. It is apparent that in order for the dynamic mode decomposition to extract pertinent flow processes, these processes must be sampled at a sufficiently high frequency. The lower bound for this sampling frequency is given by the Nyquist criterion which states that only processes can be identified that are sampled with at least twice their inherent frequency. Experience however has shown that a sampling frequency of about three times the Nyquist cutoff produces satisfactory results. By tuning the sampling frequency of our data sequence, we are able to concentrate on fast or slow processes. For slow sampling frequencies, fast processes will be sampled at near-random amplitudes and thus interpreted as incoherent noise. For high sampling rates, slow processes will be quasi-steady and be reflected in the mean mode (with no growth/decay nor any oscillatory component) or a very slow drift mode. If the sampling frequency is tuned to the frequency of an oscillatory flow, the mapping from period to period will identify the Floquet matrix whose eigenvalues represent the Floquet multipliers.

3. Applications and results

The above algorithm for extracting dynamic information from numerically generated or experimentally measured data will now be illustrated by considering three problems: the flow over a square cavity, the wake behind a flexible membrane and the instabilities observed in experiments of a jet passing between two cylinders. The first example is based on data from numerical simulations and is intended to further validate the method and demonstrate specific aspects of the dynamic mode decomposition, such as convergence behaviour and subdomain analysis. The second and third examples rely on experimental data and have been included to illustrate the dynamic mode decomposition on typical set-ups from experimental fluid dynamics, such as high-speed flow visualizations and time-resolved PIV measurements. A spatial DMD analysis is demonstrated on the second example, and a comparison with results from a POD analysis is included in the third example.

3.1. *Example 1: flow over a square cavity*

First, we consider the flow over a square cavity (see figure 4*a*). The computational domain consists of the rectangle $(x, y) \in [0, 1] \times [-1, 1]$ whose lower half represents the cavity. On the left side of the upper half, a uniform flow enters the domain; on the right side of the upper half, outflow boundary conditions are imposed. At sufficiently large Reynolds numbers this flow exhibits self-sustained oscillations of the shear layer that forms on top of the cavity (Sipp & Lebedev 2007; Barbagallo, Sipp & Schmid 2009). The ensuing unsteadiness also has an influence on the vortical flow inside the cavity and drives the shear layer that forms and detaches near the right wall of the cavity. This type of flow configuration is often taken as a model of the unsteady behaviour that is observed in separation bubbles.

Conceptually, we proceed as in the previous example by sampling a two-dimensional linearized Navier–Stokes programme based on finite differences and a fractional-step method (see Schmid 2007) at equidistant time intervals. The stored two-dimensional flow fields then comprise the columnar entries in the snapshot matrix $\boldsymbol{V}_1^N$ which is then processed according to the dynamic mode decomposition algorithm. As the number of snapshots is increased we observe convergence and saturation for the algorithm. The residual for this process, based on (2.7), is given in figure 4(*b*) as a function
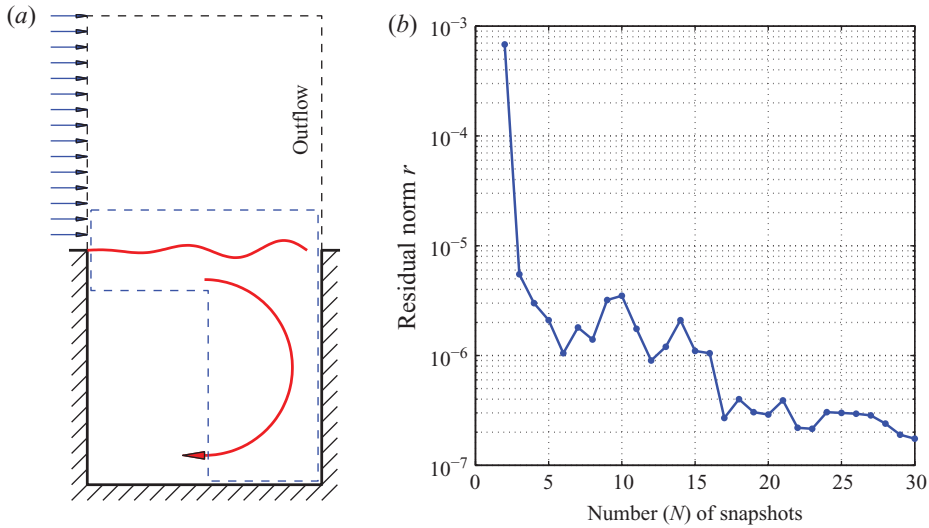
FIGURE 4. (*a*) Sketch of cavity geometry with subdomain indicated by the blue dashed line. (*b*) Residual history of the dynamic mode decomposition for cavity flow.
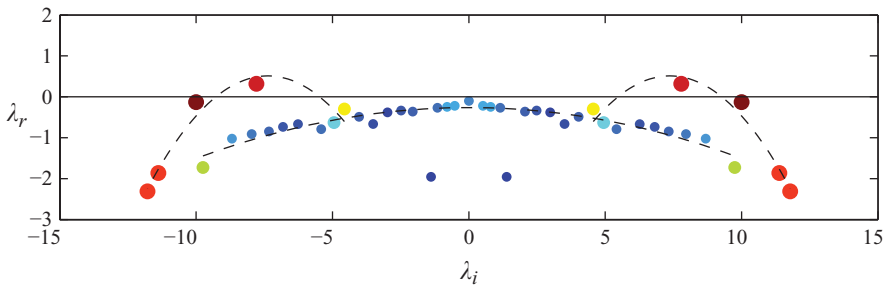

FIGURE 5. DMD spectrum for cavity flow at a Reynolds number $Re = 4500$. The colour and symbol size of the eigenvalues indicate the coherence of the associated dynamic modes.

of the number of included snapshots. Rapid convergence is observed; only about 20 snapshots are needed to determine the growth rates and phase velocities with sufficient accuracy. The extracted spectrum (i.e. the spectrum of $\tilde{\mathbf{S}}$) is displayed in figure 5. The spectrum appears symmetric with respect to the imaginary axis $\lambda_i = 0$, which is a consequence of processing real-valued data. If general complex-valued data are processed (e.g. after a Fourier transformation along a homogeneous or periodic direction), the spectrum will generally appear asymmetric with respect to $\lambda_i = 0$. For the chosen parameters ($Re = 4500$ based on the uniform inflow velocity and the cavity length), a pair of unstable modes has been identified. In addition, we observe a typical spectrum for a cavity, consisting of a parabolic branch containing the unstable eigenvalues (the unstable branches) and a parabolic branch describing the dynamics inside the cavity (the stable branch). This type of spectrum should be compared with spectra from a global stability analysis for similar flow configurations (see e.g. Åkervik *et al.* 2007 for a shallow cavity, and Sipp & Lebedev 2007 and Barbagallo *et al.* 2009 for a square cavity). The symbol size and colouring of the eigenvalues in figure 5 indicates a coherence measure of the associated modes and is intended to separate large-scale energetic structures (in red) from smaller-scale less-energetic
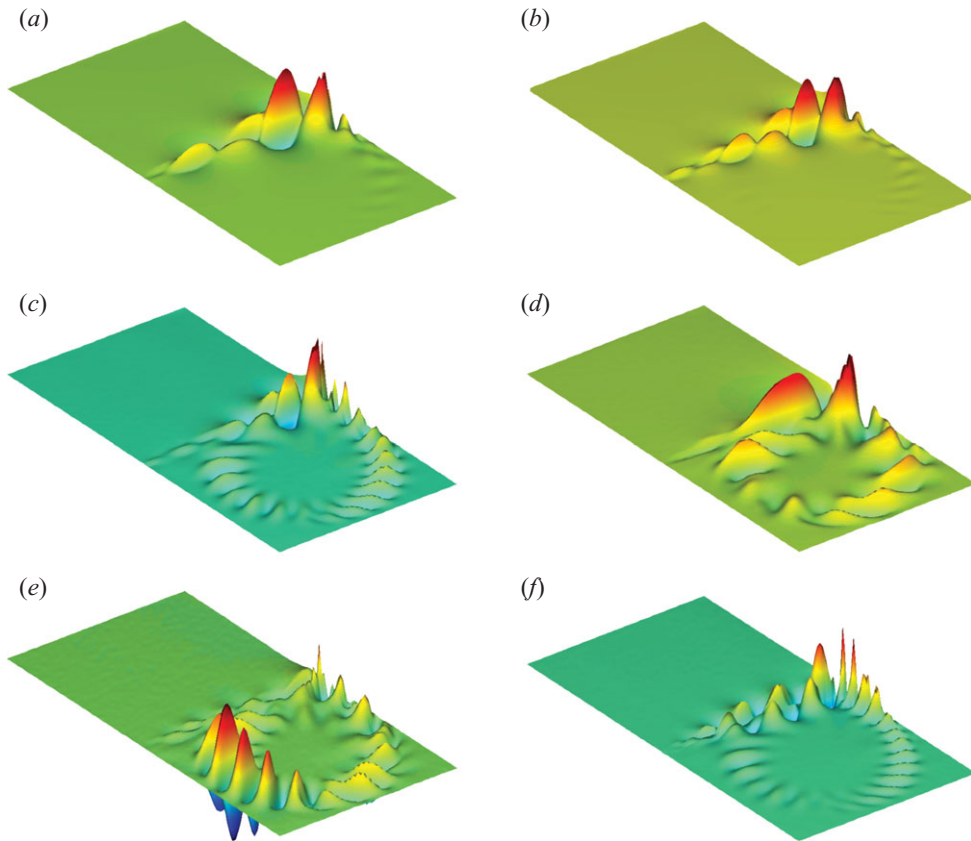
FIGURE 6. Representative dynamic modes, visualized by the streamwise velocity component, for flow over a cavity at $Re = 4500$. (a) Most unstable dynamic mode, (b–d) dynamic mode from the unstable branch, (e, f) dynamic modes from the stable branch. Because data from linearized Navier–Stokes simulations have been processes, the dynamic modes are equivalent to global modes.

structures. The criterion is given by a projection of a specific dynamic modes Φ_i onto the POD basis $\boldsymbol{U}$, computed from the data sequence $\boldsymbol{V}_1^{N-1}$; the modulus of the coefficients of this projection measures the presence of various POD modes and thus gives a measure of coherence. It is important to realize, however, that modes with a moderate to small projection onto a POD basis (blue symbols) can still play a significant dynamic role within the snapshot sequence.

Representative dynamic modes are displayed in figure 6 using the streamwise velocity component; their respective eigenvalues are given in table 1 (second and third columns). The unstable mode (figure 6a) is clearly located in the shear layer of the flow and shows the characteristic streamwise wavelength of the observed instability. Other modes from the unstable branch (figure 6b–d) have significant components in the shear layer, but also show features inside the cavity. These features are related to the instability of the shear layer detaching from the right edge of the cavity. Dynamic modes from the stable branch (figure 6e,f) contain similar characteristics: vortical structures coincidental with the mean shear layer on top of the cavity and features linked to the vortex inside the cavity. Modes from the stable branch show increasingly more small-scale features inside the cavity, as the frequency λ_i increases, which is consistent with observations of Barbagallo *et al.* (2009).

Mode	λ_r	λ_i	λ_r	λ_i
(a)	0.31705	7.79004	0.31705	7.79004
(b)	−0.13154	10.00141	−0.12997	10.00307
(c)	−1.86194	11.38202	−1.32688	11.81210
(d)	−0.29688	4.56283	−0.25238	4.63178
(e)	−0.66418	3.50959	−0.91428	3.81468
(f)	−1.72888	9.74711	−1.58432	9.00429

TABLE 1. Comparison of extracted eigenvalues via dynamic mode decomposition using data from the full domain (second and third columns) or from the subdomain (fourth and fifth columns) for the dynamic modes depicted in figures 6 and 7, respectively.

The composite characteristics of the dynamic modes, i.e. the separation of dominant shear layer and dominant cavity modes, suggest the demonstration of the decomposition technique by subdomains. For this reason, we generate snapshots that only take flow field information from a subset of the entire domain. In particular, we extract fluid velocity data from the L-shaped domain indicated by the dashed line in figure 4(*a*). These data are then processed as before. Table 1 lists the dominant eigenvalues from a full domain and a subdomain decomposition of the flow fields taking an identical number of snapshots. We observe that the dynamic modes that have their support within the subdomain are particularly well approximated, whereas moderate but noticeable differences occur for modes that are only captured partially by the subdomain. The dynamic modes from the subdomain analysis are given in figure 7 for the eigenvalues equivalent to the ones selected for figure 6. Very good agreement for the unstable modes can be observed; the agreement for the stable modes is still remarkable.

These results are rather encouraging and illustrate the potential of approximating the prevalent dynamics from snapshots even though the full extent of the instability mode is only partially (but sufficiently) captured by the measurements.

3.2. Example 2: flow in the wake of a flexible membrane

After having gained confidence in the capability of the DMD technique to extract coherent structures from snapshot sequences of temporally evolving data, we now apply this technique to data sequences from experimental measurements based on time-resolved PIV measurements of the wake flow at the edge of a flexible membrane. The geometry of the experimental set-up is displayed in figure 8, both as a conceptual sketch (figure 8*a*) and as a top view on the experiment (figure 8*b*). A uniform flow passes over a U-shaped thin steel frame that holds a flexible latex membrane. The wake flow of this membrane is affected by the rigid steel frame as well as the fluid–structure interaction between the membrane and the boundary layer that passes over it. Interest in this type of configuration comes from its relevance for micro air vehicle (μAV) performance and from a desire to better understand unsteady thrust and lift generation by using a flexible flapping membrane (see e.g. Schmit & Glauser 2009).

The dimensions of the latex membrane have been chosen as $32\,\text{mm} \times 66\,\text{mm}$ with a thickness of $0.15\,\text{mm}$. The upstream and lateral edges of the latex membrane have been attached to a $0.2\,\text{mm}$ thin U-shaped steel frame, and the free-stream velocity in the wind tunnel has been set to $14\,\text{m s}^{-1}$. The flow features in the wake have been captured by a time-resolved PIV system that has been placed in the plane of the steel frame at a distance of $38\,\text{mm}$ downstream of (but laterally shifted to) the trailing edge of the membrane. Data have been acquired from an interrogation area of
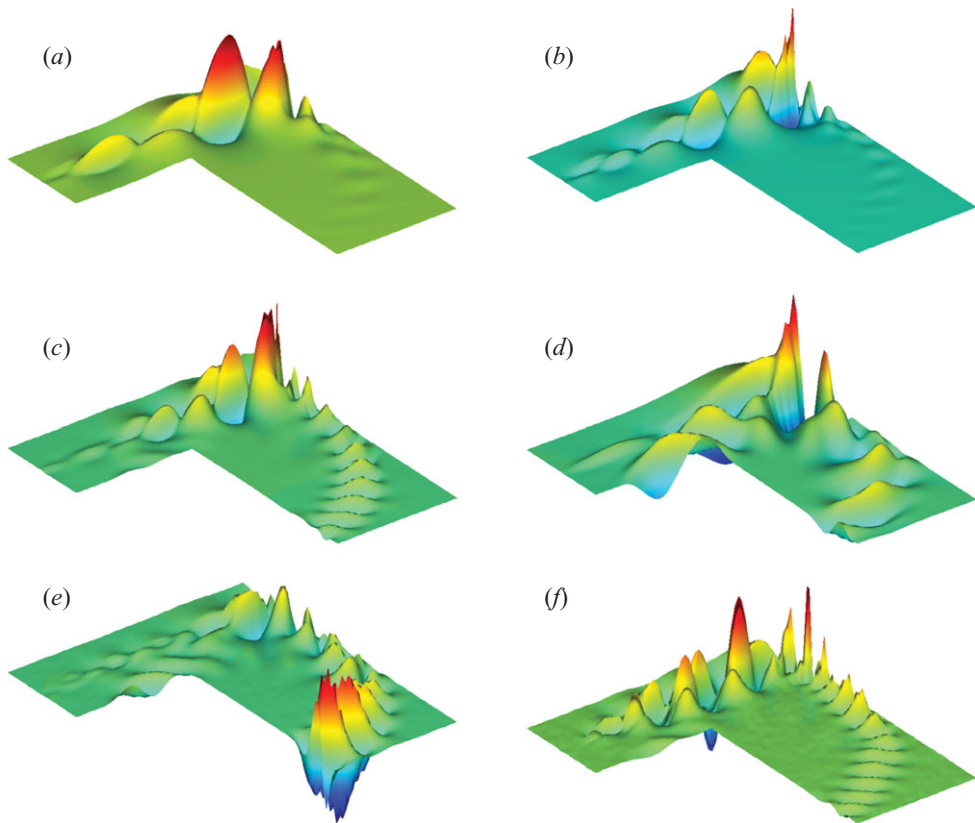
FIGURE 7. Representative dynamic modes for flow over a cavity at $Re = 4500$ sampled only in the L-shaped domain indicated in figure 4(a). The modes are associated with the selection given in figure 6.
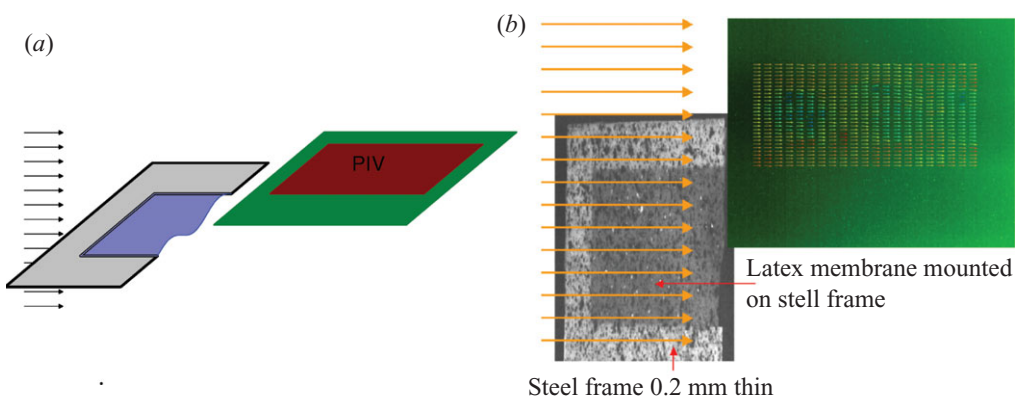

FIGURE 8. (a) Sketch of the experimental set-up. (b) Top view of the set-up for measurements in the wake of a flexible membrane; the steel frame and the PIV interrogation window are not to scale (true dimensions are given in the text). The green area on the right shows a typical snapshot from the time-resolved PIV measurements.
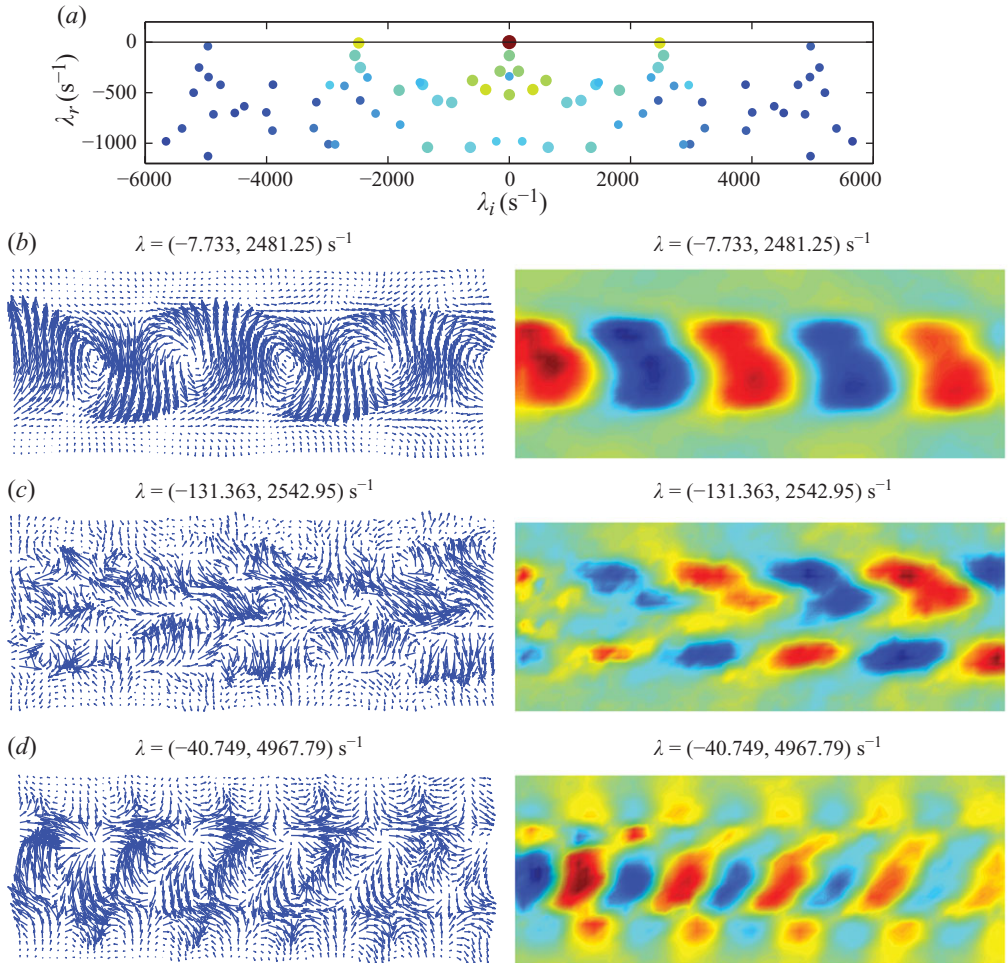
FIGURE 9. (*a*) DMD spectrum of wake flow behind a flexible membrane. (*b*–*d*) Three representative dynamic modes corresponding to dominant eigenvalues in the spectrum, visualized by vector plots (left) and contours of the transverse velocity component (right).

164 mm × 84 mm. The time-resolved camera acquired raw particle images at 2000 Hz, which were subsequently processed using an adaptive cross-correlation algorithm. The initial interrogation area size has been resolved by 64 px × 64 px. In a second and final step, measurements have been taken with a resolution of 32 px × 32 px and a 50 % overlap. An intermediate local mean validation used a 3 × 3 neighbourhood. Included in figure 8 is a typical snapshot of the flow (green area) which appears rather uniform, even though small organized and coherent fluctuations about this uniform flow are already discernible by inspection alone.

The time-resolved two-component velocity measurements on the 64 × 64 grid have then been processed to form a snapshot sequence capturing the flow evolution over a characteristic time interval. A total of 500 snapshots have been processed. A low-dimensional mapping $\tilde{\mathbf{S}}$ between these snapshots (assumed constant over the sampling interval) has then been computed and further processed to extract a DMD spectrum together with the dynamic modes of the flow. The spectrum extracted from the data is shown in figure 9(*a*), where the eigenvalues μ_i of $\tilde{\mathbf{S}}$ have been mapped according to $\lambda_i = \log(\mu_i)/\Delta t$ with $\Delta t = 5 \times 10^{-4}$ s.

Besides the time-invariant mode at the origin of the spectrum, it identifies two dominant frequencies for the flow with the higher frequency associated with smaller scales. The corresponding dynamic modes, visualized by vector plots and by contours of the streamwise velocity component, are displayed in figure 9(*b–d*) for three of the eigenvalues. The most dominant and coherent dynamic mode is characterized by large-scale vortical structures with approximately two and a half oscillatory periods over the streamwise extent of the interrogation window. Nearly the same scaling is evident from the second dynamic mode. The third dynamic mode plotted in figure 9 corresponds to the larger-frequency eigenvalue $\lambda = (-40.749, 4967.79)\,\text{s}^{-1}$ which is nearly double the frequency of the previous two modes. The spatial scale associated with this mode is correspondingly smaller. In fact, five periods are observed in the streamwise direction. In addition, a monotonic decay of the amplitude of this pattern in the streamwise direction is clearly visible.

These observations motivate the treatment of the spatial problem based on the data we have gathered by PIV. As mentioned earlier, because a low-dimensional representation of the mapping from snapshot to snapshot is one of the outputs of the dynamic mode decomposition, it is inconsequential to the algorithm whether the snapshots are aligned in time or in space. All that is required is the extraction of $(t - y)$-planes at the given streamwise x-locations. Note, however, that the spatial resolution in the x-direction has to be sufficient to resolve and detect coherent structures in space. Processes that progress faster in space than the spatial Nyquist wavenumber will not be captured by the method.

The spatial spectrum for flow in the wake of a flexible membrane is shown in figure 10 where a streamwise distance of $\Delta x = 2.5\,\text{mm}$ between the snapshots has been used. It displays one dominant mode, besides the eigenvalue at the origin. This mode has a spatial eigenvalue of $\alpha \approx (-0.0035, 0.09486)\,\text{mm}^{-1}$. The imaginary part of this eigenvalue describes the harmonic structure in the streamwise direction. Over the streamwise extent of the interrogation window (164 mm), this spatial wavenumber corresponds to a structure with 2.47 periods. Furthermore, an estimate of the temporal frequency from the oscillatory pattern shown in the corresponding subfigure in figure 10 reveals a value of about $f \approx 2356\,\text{Hz}$. These values, i.e. the spatial wavenumber and temporal frequency, can be linked to the ones that the temporal DMD analysis produced: a dominant structure with a frequency of 2481 Hz and a spatially harmonic shape with 2.5 periods inside the interrogation window (see figure 9). In addition, the noticeable slow decay in amplitude of the dominant temporal dynamic mode (shown in figure 9 for $\lambda = (-7.733, 2481.25)\,\text{s}^{-1}$) can be determined by fitting an exponential function to the peaks of the streamwise velocity component; the resulting decay rate is $\gamma = -0.003355\,\text{mm}^{-1}$, in very good agreement with the real part of the identified spatial eigenvalue $\alpha = (-0.00351, 0.09486)\,\text{mm}^{-1}$. A correspondence between the two types of analysis has thus been established that gives further confidence in the capability of the dynamic mode decomposition to capture the relevant processes, be it in the temporal or spatial framework. A higher temporal frequency is noticeable for a higher spatial mode (with $\alpha_i = 0.19851\,\text{mm}^{-1}$), but the limited resolution in the streamwise coordinate direction translates into a poor representation of the time-harmonic component in the associated spatial DMD mode.

3.3. *Example 3: experiment of a jet between two cylinders*

The final example is again based on data from experimental measurements using time-resolved PIV. In this case, the flow through an array of cylinders is studied. This
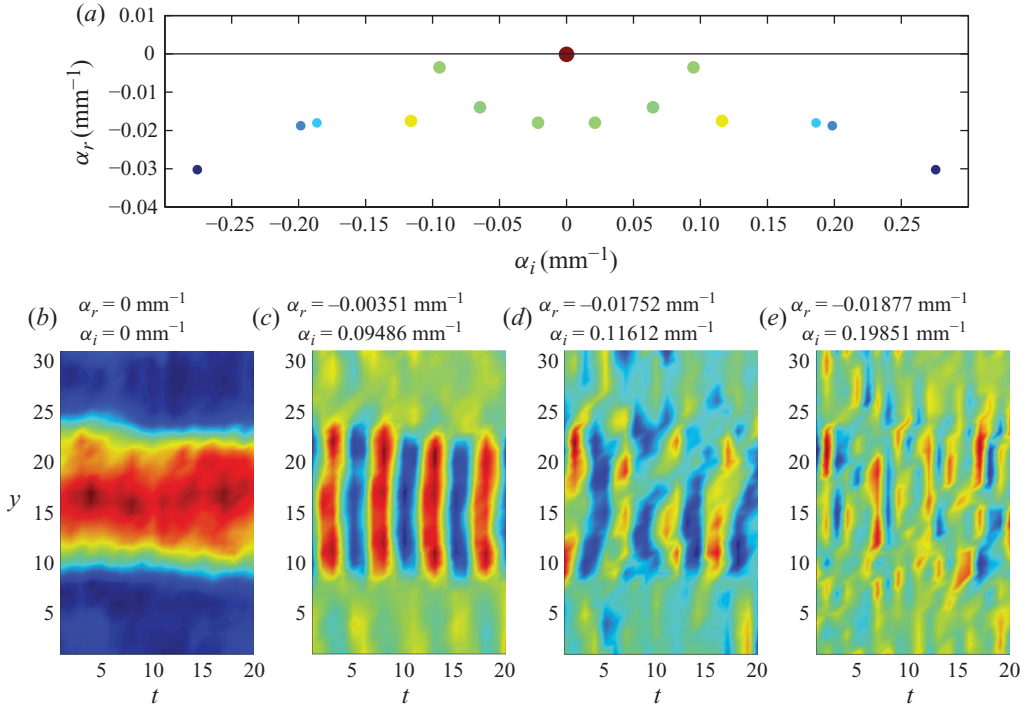
FIGURE 10. (*a*) Spatial DMD spectrum and (*b–e*) spatial dynamic modes for flow in the wake of a flexible membrane. The dynamic modes are visualized by contours of the streamwise velocity component.

flow configuration arises in many industrial settings and is of particular interest in fluid devices for energy conversion, e.g. heat exchangers.

In our setting, the flow of a jet passing between two cylinders is investigated. A sketch of the geometric set-up is displayed in figure 11(*a*), together with the interrogation window where PIV measurements of the flow are taken. An interrogation window of 40.36 mm in the streamwise and 32.08 mm in the transverse directions has been used, which has been resolved by a 63×79 measurement grid on which two velocity components have been recorded. The flow has been sampled in a time-resolved manner with 4 ms between two consecutive PIV measurements. The jet has a width of 10.7 mm and a mean velocity of 0.663 m s^{-1}. The cylinders have a diameter of 12 mm, and the Reynolds number, based on the volume flux velocity (with $\dot{Q} = 18$ m^3 h^{-1}) and the cylinder diameter, is $Re = 3000$. In the same figure, five snapshots of raw PIV data are depicted to give a first impression of the flow field. The strong jet, forced through the gap of the two cylinders, interacts with the cylinder wakes and produces a complex flow pattern consisting of intermittent coherent vortices.

An identification of relevant structures for this type of flows has traditionally relied on the POD which forms the (temporally averaged) spatial correlation tensor which is in turn diagonalized by an eigenvalue decomposition. The latter step decorrelates the extracted structures and provides a hierarchy of coherent fluid elements ranked by their energy (or autocorrelation) content. Alternative to forming the correlation tensor, a singular value decomposition of the snapshot sequence can be taken. The
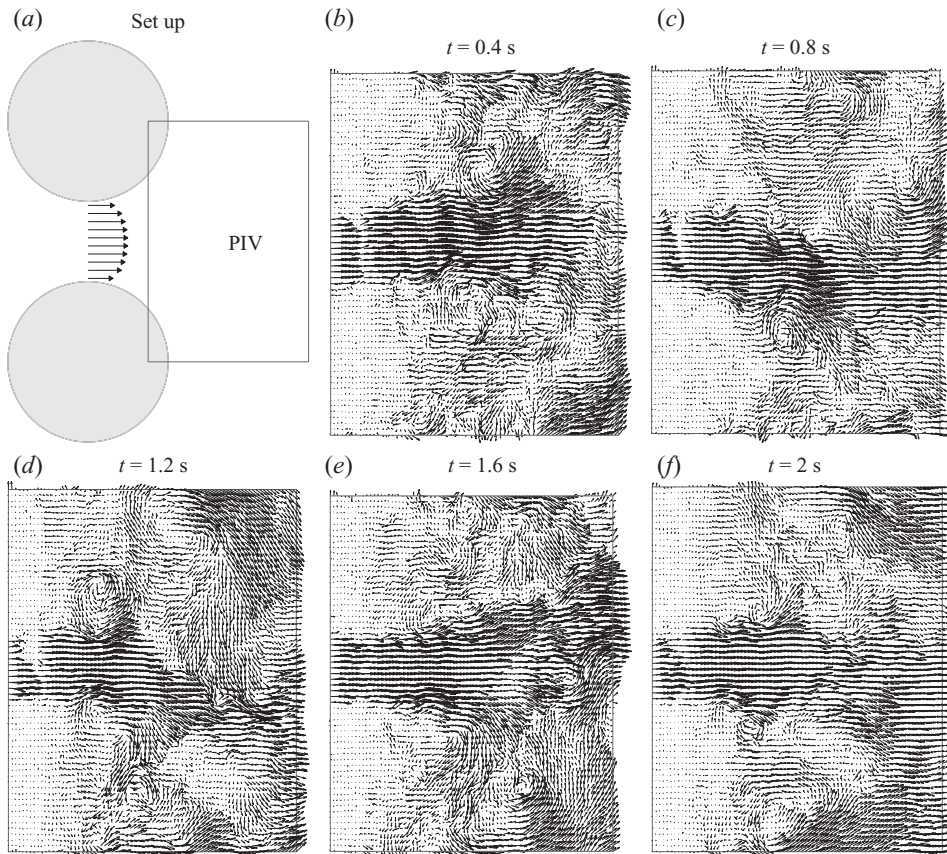
FIGURE 11. Sketch of the experimental set-up (*a*) and representative raw PIV data from time-resolved measurements (*b*–*f*). Data courtesy of EdF.

right singular vectors then contain the POD modes, while the singular values produce the energy content of each associated structure.

For the above experimental data and for comparison purposes, a POD analysis using 500 snapshots has been performed, which is displayed in figure 12. The singular values, shown in figure 12(*a*), show a rather weak decay after the first initial drop. The first POD mode, corresponding to the largest singular value, reproduces the mean flow and displays a strong jet, together with two weak lateral vortices and an outer entrainment flow. The second POD mode, already with a substantially reduced energy content, shows an antisymmetric pair of vortices located near the lateral shear layers of the jet. The third and fourth POD modes, still (anti)symmetric with respect to the jet axis, contain four and two major vortices, respectively. Starting with the fifth POD mode, at 15 % of the mean energy, first signs of an asymmetric structure appear. Despite this breakdown of the flow into energetic coherent structures, no information can be gained about the 'dynamics' of these structures. This is a consequence of the time-averaging process to obtain the spatial correlation tensor. Techniques exist, however, to recover phase information from the data set (see Lumley 1970).

In contrast, the dynamic mode decomposition does not extract flow information based on the energy content of the identified flow structures; instead, it approximates the temporal 'dynamics' of the flow by a linear evolution operator – expressed in the
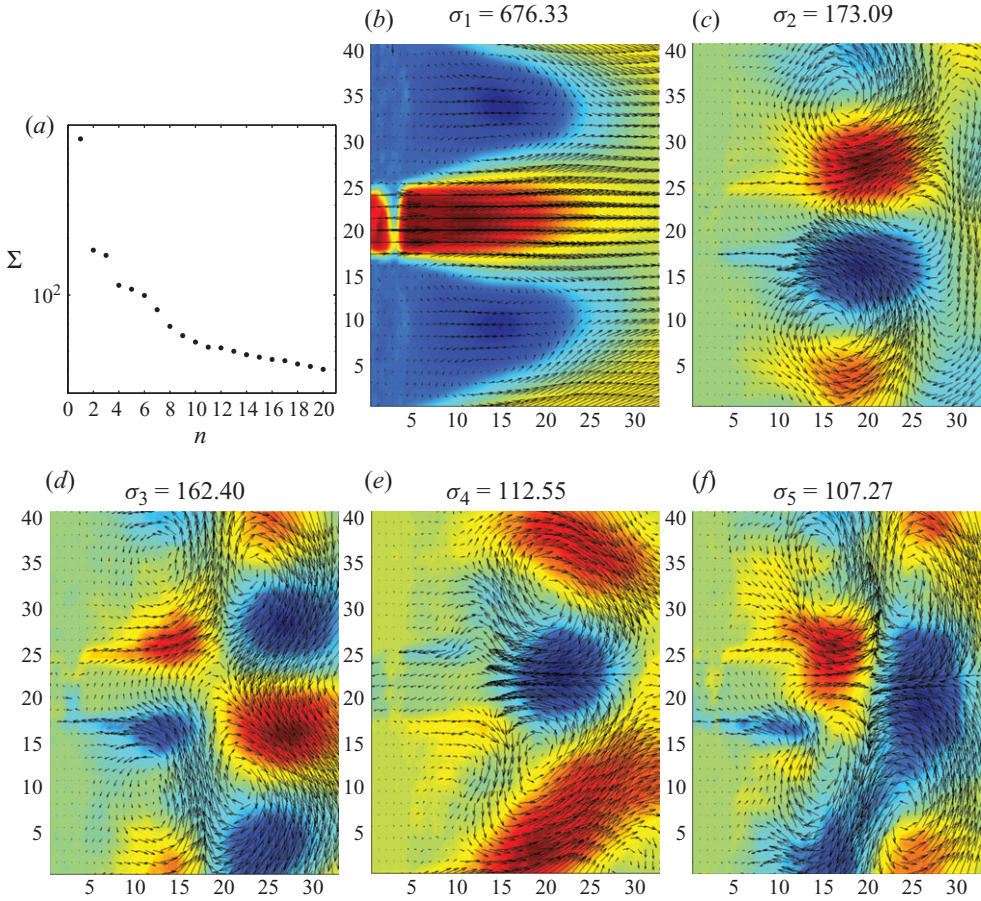
FIGURE 12. (*a*) Singular values of the snapshot sequence for flow through an array of cylinders; (*b–f*) POD modes. The background colour contours visualize the streamwise velocity component.

snapshot basis – whose dominant eigenfunctions form a set of dynamically relevant modal structures. Temporal information is thus contained. Results from this type of analysis (using $N = 200$ snapshots) are shown in figure 13. The temporal spectrum shows eigenvalues that approximately fall on parabolic curves, a feature that is often observed in global stability analyses and that reflects the presence of advective and diffusive flow phenomena. In a local stability analysis, each parabolic arc would correspond to a specific characteristic disturbance wavelength in the streamwise direction; in a global stability analysis, no such separation of scales in the streamwise direction is assumed, and the parabolic arcs for each scale appear simultaneously. The eigenvalues near the apex of these parabolic arcs correspond to the most dominant structure of the respective family of coherent elements. They are displayed, for the three arcs indicated in the spectrum, in figure 13 together with the mean flow. It is interesting to note that the identified mean flow is not strictly symmetric with respect to the jet axis. This asymmetry is even more visible in the first (and most dominant) dynamic mode, associated with the eigenvalue $\lambda = (0.210, 50.254)\,\mathrm{s}^{-1}$. Two strong vortices are present on either side of the jet, similar to the second POD mode, but the upper side of the jet has been identified as a region of increased vortical activity
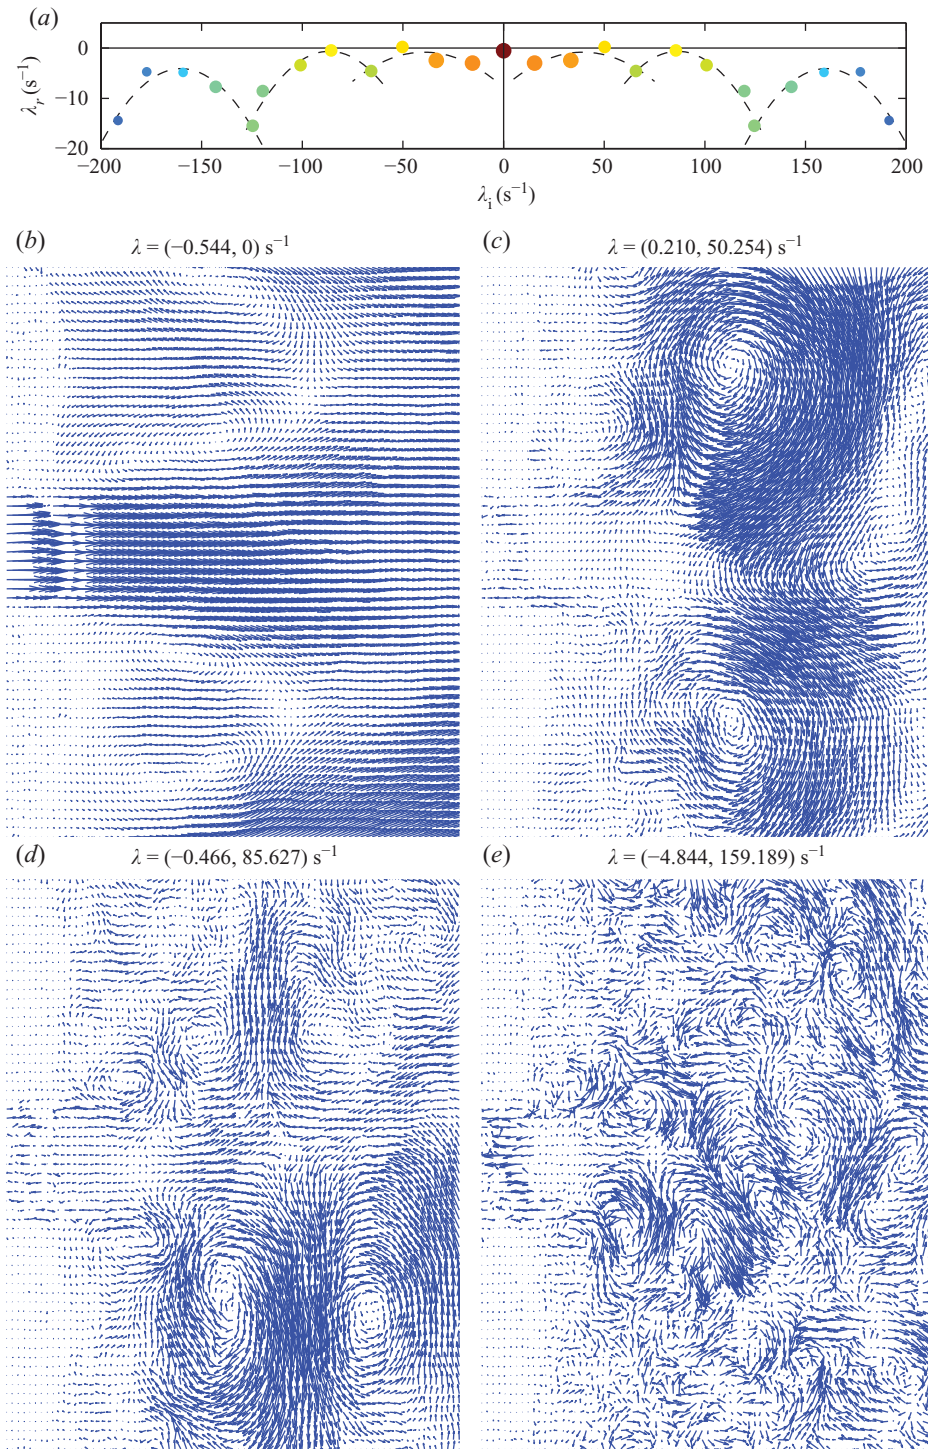
FIGURE 13. (*a*) DMD spectrum and (*b–e*) four dynamic modes representing the mean flow and three dominant modes from the apex of the parabolic arcs indicated in the spectrum.

(which can be confirmed by inspecting the raw PIV fields). The preference of one side of the jet over the other is even more pronounced in the dominant dynamic mode of the second parabolic arc ($\lambda = (-0.466, 85.627)\,\text{s}^{-1}$). In this case, a strong side jet caused by a pair of counter-rotating vortices is observable at the lower side of the jet. The upper side also contains a side jet, however, with a substantially weaker amplitude. A similar flow feature appears in the fifth POD mode (see figure 12), but the marked difference between both sides of the jet appears to have been lost in the averaging process, in addition to any information about its dynamic behaviour. The dominant dynamic mode of the third parabolic arc (with $\lambda = (-4.844, 159.189)\,\text{s}^{-1}$) displays vortical structures of a smaller scale, consistent with its higher associated frequency and decay rate. Nevertheless, coherent oblique side jets are discernible with a preferred appearance on the lower jet shear layer.

The dynamic mode decomposition appears capable of extracting dominant flow features from the snapshot basis. It does so by approximating the linear mapping between the snapshots and subsequently detecting the pertinent frequencies. In contrast, POD uses a second-order statistics of the flow fields and produces a hierarchy of coherent structures that diagonalize their correlation tensor. One can say that, while POD concentrates on a representation based on spatial orthogonality, DMD focuses on a representation based on temporal orthogonality (frequencies). For the above case of the flow between two cylinders, the latter decomposition produces a more accurate and complete description of the 'dynamic' flow behaviour for the experimentalist, capable of identifying dynamic features of the flow directly, rather than through its second-order statistics.

4. Summary and conclusions

The search for physical mechanisms underlying fluid flow relies on the decomposition of flow fields into coherent structures and flow patterns whose dynamics provide a more compact and instructive manner of describing the fluid process. For numerical simulations the Arnoldi method and various other decomposition techniques provide the tools to extract this desired information. For experimental data, on the other hand, no such decomposition exists that could provide temporal dynamic characteristics together with spatially coherent structures.

The dynamic mode decomposition offers a tool for extracting dynamic information from a sequence of uniformly sampled flow measurements. The resulting modes represent the relevant flow structures that contribute most to the overall 'evolution' captured in the measurement sequence. The technique is flexible enough to equally deal with simple flow visualizations (as demonstrated by decomposing Schlieren images in Schmid *et al.* 2010) and with time-resolved PIV measurements. Subdomains within the flow configurations can be analysed in isolation to focus on particular flow features and instability mechanisms. This is of particular advantage for flows containing a multitude of instability mechanisms or multiphysics phenomena (e.g. shear instabilities and acoustic radiation). In fact, the flow fields $\boldsymbol{v}_i$ may even contain data from different regions and/or contain different flow variables. For example, the simultaneous processing of PIV velocity measurements and of acoustic pressure signals from an array of microphones is conceivable.

Spatially evolving flows are as simple to decompose as temporally evolving ones. A simple reorganization of the measured data and an alignment of the snapshot basis along a spatial, rather than temporal, dimension yield a snapshot-to-snapshot mapping that contains information about the spatial evolution of coherent structures.

This type of analysis is of particular importance, e.g. for flows shedding from blunt bodies or localized roughness elements.

The dynamic mode decomposition is different from the proper orthogonal decomposition where the former attempts to represent a data sequence by orthogonalizing it in time (i.e. isolating distinct frequencies in the data), while the latter attempts a decomposition based on orthogonality in space. Furthermore, the dynamic mode decomposition applies directly to the data, while a POD analysis processes second-order statistics of the data.

When applied to a linear process, the dynamic modes recover the global stability modes of the flow. When decomposing a nonlinear process, the eigenvalues of the matrix S (or $\tilde{S}$) approach the unit circle in the limit of an infinite (or a sufficiently large) number of snapshots (see e.g. Rowley *et al.* 2009). In this limit, the dynamic mode decomposition identifies the dominant frequencies and their associated spatial structures. In this manner, even turbulent flows can be analysed as to the existence of cyclical coherent events, following the suggestions of Hussain (1986).

The dynamic mode decomposition has been illustrated on a variety of examples ranging from plane Poiseuille flow (for validation purposes) to flow over a cavity, from wake flow behind a flexible membrane measured by time-resolved PIV to flow of a jet passing between two cylinders. During these tests, and many more not included in this paper, the dynamic mode decomposition has proven itself as a robust and reliable algorithm to extract spatio-temporal coherent structures from the data. Further potential applications of this decomposition include, among others, the recovery of flow fields from gappy data sequences, the model reduction of high degrees-of-freedom systems, and an image-based quantitative analysis when combined with optical flow techniques. In closing, it is hoped that this type of decomposition will join the current arsenal of tools for experimentalists (and computational fluid dynamicists) and will help them in their efforts to quantify fluid-dynamical processes by their dominant spatio-temporal behaviour.

REFERENCES

ÅKERVIK, E., HŒPFFNER, J., EHRENSTEIN, U. & HENNINGSON, D. S. 2007 Optimal growth, model reduction and control in a separated boundary-layer flow using global modes. *J. Fluid Mech.* **579**, 305–314.

AUBRY, N. 1991 On the hidden beauty of the proper orthogonal decomposition. *Theor. Comp. Fluid Dyn.* **2**, 339–352.

BARBAGALLO, A., SIPP, D. & SCHMID, P. J. 2009 Closed-loop control of an open cavity flow using reduced-order models. *J. Fluid Mech.* **641**, 1–50.

BERKOOZ, G., HOLMES, P. & LUMLEY, J. L. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25**, 539–575.

BONNET, J.-P., COLE, D. R., DELVILLE, J., GLAUSER, M. N. & UKEILEY, L. S. 1994 Stochastic estimation and proper orthogonal decomposition: complementary techniques for identifying structure. *Exp. Fluids* **17**, 307–314.

DELSOLE, T. & HOU, A. Y. 1999 Empirical stochastic models for the dominant climate statistics of a general circulation model. *J. Atmos. Sci.* **56**, 3436–3456.

Edwards, W. S., Tuckerman, L. S., Friesner, R. A. & Sorensen, D. C. 1994 Krylov methods for the incompressible Navier–Stokes equations. *J. Comput. Phys.* **110**, 82–102.

Greenbaum, A. 1997 *Iterative Methods for Solving Linear Systems*. SIAM.

Hasselmann, K. 1988 POPs and PIPs: the reduction of complex dynamical systems using principal oscillations and interaction patterns. *J. Geophys. Res.* **93**, 10975–10988.

Hemon, P. & Santi, F. 2007 Simulation of a spatially correlated turbulent velocity field using biorthogonal decomposition. *J. Wind Engng Ind. Aerodyn.* **95**, 21–29.

Herzog, S. 1986 The large scale structure in the near-wall region of turbulent pipe flow. PhD dissertation, Department Mechanical Engineering, Cornell University.

Hussain, A. K. M. F. 1986 Coherent structures and turbulence. *J. Fluid Mech.* **173**, 303–356.

Lasota, A. & Mackey, M. C. 1994 *Chaos, Fractals and Noise: Stochastic Aspects of Dynamics*. Springer.

Lehoucq, R. B. & Scott, J. A. 1997 Implicitly restarted Arnoldi methods and subspace iteration. *SIAM J. Matrix Anal. Appl.* **23**, 551–562.

Lumley, J. L. 1970 *Stochastic Tools in Turbulence*. Academic Press.

Mezić, I. 2005 Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.* **41**, 309–325.

Noack, B. R., Afanasiev, K., Morzynski, M., Tadmor, G. & Thiele, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.

Noack, B. R., Schlegel, M., Ahlborn, B., Mutschke, G., Morzynski, M., Comte, P. & Tadmor, G. 2008 A finite-time thermodynamics formalism for unsteady flows. *J. Non-Equilib. Thermodyn.* **33**, 103–148.

Orszag, S. A. 1971 Accurate solution of the Orr–Sommerfeld stability equation. *J. Fluid Mech.* **50**, 689–703.

Penland, C. & Magoriam, T. 1993 Prediction of Niño 3 sea-surface temperatures using linear inverse modelling. *J. Climate* **6**, 1067–1076.

Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P. & Henningson, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.

Ruhe, A. 1984 Rational Krylov sequence methods for eigenvalue computation. *Linear Algebr Appl.* **58**, 279–316.

Schmid, P. J. 2007 Transition and transition control in a square cavity. In *Advances in Turbulence XI* (ed. J. M. L. M. Palma & A. Silva Lopez), pp. 562–569. Springer Verlag.

Schmid, P. J. & Henningson, D. S. 2001 *Stability and Transition in Shear Flows*. Springer.

Schmid, P. J., Li, L., Juniper, M. P. & Pust, O. 2010 Applications of the dynamic mode decomposition. *Theor. Comput. Fluid Dyn.* (in press).

Schmid, P. J. & Sesterhenn, J. L. 2008 Dynamic mode decomposition of numerical and experimental data. In *Bull. Amer. Phys. Soc., 61st APS meeting*, p. 208. San Antonio.

Schmit, R. F. & Glauser, M. N. 2009 Use of low-dimensional methods for wake flowfield estimation from dynamic strain. *AIAA J.* **43** (5), 1133–1136.

Sipp, D. & Lebedev, A. 2007 Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows. *J. Fluid Mech.* **593**, 333–358.

Sirovich, L. 1987 Turbulence and the dynamics of coherent structures. *Q. Appl. Math.* **45**, 561–590.

von Storch, H., Bürger, G., Schnur, R. & von Storch, J. 1995 Principal oscillation pattern: a review. *J. Climate* **8**, 377–400.

Trefethen, L. N. & Bau, D. 1997 *Numerical Linear Algebra*. SIAM.