J. Fluid Mech. (2024), *vol*. 1001, A48, doi:10.1017/jfm.2024.1097

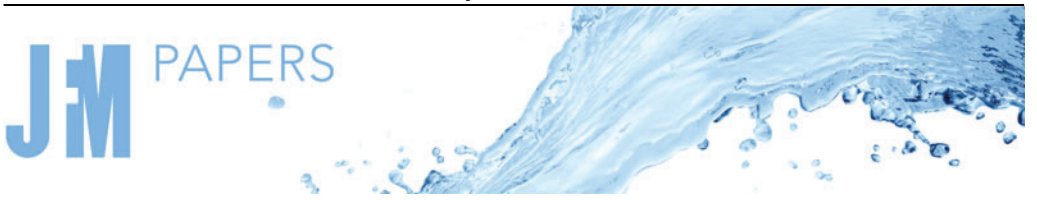

Dynamics-preserving compression for modal flow analysis

Anton Glazkov[1] **and Peter J. Schmid**[1],†

[1]Department of Mechanical Engineering, Division of Physical Sciences and Engineering (PSE), King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia

(Received 14 November 2023; revised 29 September 2024; accepted 6 November 2024)

An efficient compression scheme for modal flow analysis is proposed and validated on data sequences of compressible flow through a linear turbomachinery blade row. The key feature of the compression scheme is a minimal, user-defined distortion of the mutual distance of any snapshot pair in phase space. Through this imposed feature, the model reduction process preserves the temporal dynamics contained in the data sequence, while still decreasing the spatial complexity. The mathematical foundation of the scheme is the fast Johnson–Lindenstrauss transformation (FJLT) which uses randomized projections and a tree-based spectral transform to accomplish the embedding of a high-dimensional data sequence into a lower-dimensional latent space. The compression scheme is coupled to a proper orthogonal decomposition and dynamic mode decomposition analysis of flow through a linear blade row. The application to a complex flow-field sequence demonstrates the efficacy of the scheme, where compression rates of two orders of magnitude are achieved, while incurring very small relative errors in the dominant temporal dynamics. This FJLT technique should be attractive to a wide range of modal analyses of large-scale and multi-physics fluid motion.

Key words: computational methods, big data

1. Introduction

As we tackle complex and large-scale fluid problems with numerical and experimental methods, the proper analysis of these flows – while increasingly challenging – becomes ever more essential for our understanding of the underlying flow physics. A key tool in these analyses are decompositions of the flow fields (Taira *et al.* 2017), either purely data-based or linked to a set of governing equations, that produce a hierarchy of coherent

† Email address for correspondence: peter.schmid@kaust.edu.sa

structures, key processes and persistent dynamics, and break down a complex flow behaviour into its simpler constituent parts. This hierarchical approach is predicated on the assumption that the observed flow fields, in their full complexity, can be well approximated by a far lower-dimensional set of mechanisms. These decompositions can thus be considered as model reduction techniques tailored to the specifics of the intrinsic flow physics. Various linear and nonlinear methods are available, with each approach aiming at a specific aspect of the flow and characterized by strengths and limitations.

Common to most techniques is a separation-of-variables approach that isolates normalized spatial structures from their temporal dynamics and amplitudes. Constraints on either the spatial or temporal elements of the chosen decomposition determine the enforced coherence or structural independence (e.g. orthogonality, decorrelation) in the reduced-order model. The amplitudes (or spectra) furnish a hierarchical ranking of the identified subprocesses and, together with the temporal dynamics, drive the evolution of the reduced state vector in latent space. Among commonly used decompositions, the proper orthogonal decomposition (POD; Lumley 1970; Sirovich 1987; Berkooz, Holmes & Lumley 1993; Holmes *et al.* 2012; Schmidt & Schmid 2019) plays a key role in extracting spatially decorrelated coherent structures responsible for a maximum of variance in terms of kinetic energy. Decompositions into pure-frequency mechanisms, such as resolvent analysis (McKeon 2017; Ribeiro, Yeh & Taira 2020; Rigas, Sipp & Colonius 2021), dynamic mode decomposition (DMD; Rowley *et al.* 2009; Schmid 2010; Kutz *et al.* 2016; Schmid 2022) or spectral POD (sPOD; Towne, Schmidt & Colonius 2018; Schmidt & Colonius 2020), have joined the spatially oriented factorizations of the flow fields and have produced new insight into key flow processes. Other decompositions, such as independent component analysis (Hyvärinen, Karhunen & Oja 2001) or linear stochastic estimation (Adrian & Moin 1988), are perhaps less prevalent but have nonetheless advanced our understanding of coherence in turbulent fluid motion. More recently, nonlinear model reduction techniques have been devised that account for nearest-neighbour or tangent-space measures in their algorithmic steps (see Lawrence (2012) for a review).

In these efforts, little attention is usually directed towards mapping dynamically coherent features in full space onto equally coherent features in latent space, particularly in the temporal part of the decompositions. As a result, the dynamical behaviour in latent space is often compromised in its ability to properly and concisely represent a particular motion in physical space. For example, it is commonly acknowledged that higher-order POD modes are associated with multi-frequential mixtures that, when interpreted in isolation, have little resemblance to observable mechanisms in the original high-dimensional space. Rather, a superposition of multiple modes is required to approximate temporally coherent (non-periodic) features in the full data sequence. Similarly, the pure-frequency decompositions (resolvent, DMD, sPOD) capture periodic flow features with distinct Strouhal numbers, but struggle with more complex and broadband time behaviour.

To compensate for this shortcoming, in particular in data-driven approaches, time-delay or Ruelle–Takens embedding has been proposed (Takens 1981). This technique maintains temporal coherence across multiple time steps by considering sequences of snapshots as state vectors rather than individual flow fields. The key idea behind processing these short, concatenated time traces is a better retention of time coherence and a better encapsulation of history effects during the model reduction process. It is aligned with the general concept of temporally coherent structures and improves on the more commonly applied mappings over only a snapshot pair, thus moving from a Markovian to a non-Markovian description. Time-delay embedding, which leads to data matrices of block Hankel type, has been

applied with some success, such as in singular spectral analysis (Golyandina, Nekrutin & Zhigljavsky 2001), Hankelized DMD (Arbabi & Mezić 2017) or HAVOK (Brunton *et al.* 2017), in the analysis of complex flows as well as the prediction of fluid processes as part of active control efforts (Budišić, Mohr & Mezić 2012; Kamb *et al.* 2020; Pan & Duraisamy 2020; Axås & Haller 2023). The technique, however, comes at significantly increased computational costs, as each new snapshot now consists of a sizeable sequence of original snapshots which, for complex flows, are already of formidable size. This escalation of input dimensions puts a marked strain on standard (linear-algebra) algorithms used to analyse the flow.

In this present paper, we propose and validate an efficient method that reduces the degrees of freedom of the incoming snapshots, but enforces a crucial constraint on the latent space dynamics. By doing so, it accomplishes two desired outcomes: a reduced latent space and a high-fidelity correspondence between the physical and latent dynamics. These seemingly opposing goals can be accomplished with a remarkable technique that embeds a sequence of fields into a lower-dimensional space while preserving (within user-specified tolerances) their mutual distances via randomized projections. In combination, it thus acts as a constrained, dynamics-preserving model reduction strategy.

The rise of randomized methods, in particular in numerical linear algebra (Halko, Martinsson & Tropp 2011; Martinsson & Tropp 2020), has also reached quantitative flow analysis in the form of highly efficient algorithms that extract coherent structures from the response of the fluid system to a random forcing (Ribeiro *et al.* 2020). These techniques, which also go by the name of matrix sketching, play a key role in applying standard analysis tools to large-scale and fully three-dimensional flows, and do so in a computationally efficient manner. While the present method falls within the category of randomized methods, it crucially distinguishes itself by its dynamics-preserving property. None of the previous sketching techniques comes with the guarantees of minimal distance distortion among any of the processed snapshots, and therefore maps the full dynamics into latent space with a significant amount of distortion and a potentially suboptimal representation.

We first sketch the theoretical background for the dynamics-preserving model reduction technique based on randomized projections. This technique is then attached to the procedural steps of a modal analysis, based on POD and DMD. An application to flow fields from simulations of compressible flow through a linear blade cascade will demonstrate the efficacy of the new compression technique, showing that the key mechanisms contained in the snapshot sequence can be faithfully extracted directly from the compressed sequence at significantly reduced computational cost. A discussion of future research directions and open challenges concludes this study.

2. Theoretical background

We start with a data sequence $\boldsymbol{Q} = \{\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots, \boldsymbol{q}_n\}$ consisting of n snapshots of observables where each member $\boldsymbol{q}_i \in \mathbb{R}^d$ of the set represents a high-dimensional state vector with d degrees of freedom and sampled from numerical simulations or physical experiments. Our task then consists of designing a mapping $\boldsymbol{\Phi}$ that reduces each high-dimensional $\boldsymbol{q}_i$ in the data sequence to a lower-dimensional equivalent $\boldsymbol{y}_i = \boldsymbol{\Phi}\boldsymbol{q}_i \in \mathbb{R}^k$ with $k \ll d$ while observing a given constraint. This mapping $\boldsymbol{\Phi}$ will be linear and, hence, can be represented as a $k \times d$ matrix. As the constraint in this mapping, we aim at preserving the dynamics encoded in an N-member sample $\boldsymbol{Q}_N \subseteq \boldsymbol{Q}$, extracted from the original data sequence $\boldsymbol{Q}$. As a measure of the dynamics, we take the mutual

pairwise Euclidean distance between snapshots taken from $\boldsymbol{Q}_N$. This choice stems from an interpretation of the data sequence as a path in d-dimensional phase space: by preserving the Euclidean distance between any, and not necessarily consecutive, pair of snapshots from any N-component subset $\boldsymbol{Q}_N$, we maintain the global relational structure of the phase-space path – and thus the dynamics contained in the data sequence $\boldsymbol{Q}$. Unfortunately, this preservation cannot be enforced exactly. Instead, we have to allow a slight distortion between the distance measure of the original and reduced data sequence. By forcing the preservation of the Euclidean distance between any N snapshots, we are in effect partitioning the d-dimensional space into k regions rather than enforcing a weaker chaining-preservation property that would arise if we were to preserve just consecutive pairs in the trajectory, and accordingly this compression scheme is agnostic to the temporal dynamics. For this reason, the spatial compression will affect the temporal evolution only through the removal of some information in the compressed fields, as the dynamics is itself enforced through the ordering of the snapshots in the data sequence formed by the columns of $\boldsymbol{Q}$. The constraint on our dynamics-aware, rather than data-aware, model reduction can thus be stated as

$$(1-\varepsilon)\underbrace{\|\boldsymbol{q}_i-\boldsymbol{q}_j\|_2}_{\substack{\textit{original}\\ \textit{distance}}} \leq \underbrace{\|\boldsymbol{\Phi}\boldsymbol{q}_i-\boldsymbol{\Phi}\boldsymbol{q}_j\|_2}_{\substack{\textit{compressed}\\ \textit{distance}}} \leq (1+\varepsilon)\underbrace{\|\boldsymbol{q}_i-\boldsymbol{q}_j\|_2}_{\substack{\textit{original}\\ \textit{distance}}} \quad \text{for all } i,j\in[1,\ldots,N], \tag{2.1}$$

where $\varepsilon \ll 1$ represents the distortion factor of the mapping $\boldsymbol{\Phi}$. The existence of such a compressive linear mapping $\boldsymbol{\Phi}$ has been established via the seminal Johnson–Lindenstrauss (JL) lemma (Johnson & Lindenstrauss 1984). Besides its impact on many fields, this lemma has substantiated nearly orthogonal (approximate rotation) matrices, when acting on sparse vectors, which have become a key concept of compressed sensing and sparse signal recovery (Candès & Tao 2006). Also, the above bounds (2.1) are not universally valid, but are implied statistically. They thus come with a failure probability δ which is bounded from above by one-third.

The JL lemma and the above bounds (2.1) are formulated in the Euclidean norm, and it is a valid concern to call into question the physical nature of this choice for both the original and reduced snapshot pairs. While weighted norms are certainly conceivable, we will proceed with the native formulation (Johnson & Lindenstrauss 1984) and base our phase-space embeddings on this generic measure.

The practical realization of the mapping $\boldsymbol{\Phi}$ invokes embeddings by random projections (Achlioptas 2003), which is a path also followed by sketching techniques. The underlying idea is that probing a matrix by random vectors can reveal its main features, if the matrix, either from a data sequence or from discretized governing equations, is information-sparse when expressed in a proper basis. The research that followed the establishment of a JL-based, dimensionality-reducing mapping concentrated on two aspects: (1) the transformation of an arbitrary input signal/vector into a vector that is amenable to random projections and (2) the sparsification of the random projections themselves. While the first advance decreases the failure probability for (2.1), the second push increases the efficiency and compression rate of the embedding.

In our context, the parameter N quantifies the mutual linkage across snapshots. While a minimum of $N=2$ cross-connections are necessary, a larger value of N will lead to more accurate mappings and a reduced failure probability, at the expense of a reduced compression rate. This lower failure probability stems from the fact that for N connections, we have $\binom{N}{2}$ pairings, any one of which will fail with a lower probability to cause an overall

violation of the bound (2.1). We therefore establish a link between N and δ of the form $\delta = \frac{1}{3}\binom{N}{2}^{-1} \sim 1/N^2$. It is also important to recognize that the linkage parameter N in the fast JL transformation (FJLT) plays the role of the embedding dimension for a time-delay (Ruelle–Takens) embedding.

In the development of a practical algorithm to determine the mapping $\boldsymbol{\Phi}$, Ailon & Chazelle (2009, 2010), and later Ailon & Liberty (2013), improve on an earlier attempt by Achlioptas (2003) and propose a three-component transformation with a mapping of the form

$$\boldsymbol{\Phi} = \boldsymbol{PHD}. \tag{2.2}$$

In this formulation, the matrices $\boldsymbol{P}$ and $\boldsymbol{D}$ are random, while $\boldsymbol{H}$ is deterministic. Both $\boldsymbol{H}$ and $\boldsymbol{D}$ are $d \times d$ matrices responsible for the preconditioning of the input signal; $\boldsymbol{P}$ is a $k \times d$ sketching matrix responsible for the final compression. Following Ailon & Chazelle (2009), we have (assuming, without loss of generality, that d is a power of two)

$$\boldsymbol{D} = \mathrm{diag}\{\pm 1, \ldots, \pm 1\}, \quad \boldsymbol{H}_m = \boldsymbol{H}_1 \otimes \boldsymbol{H}_{m-1}, \quad \boldsymbol{H}_1 = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \tag{2.3}$$

with $\otimes$ representing the Kronecker product and m ranging from 2 to $\log_2 d$. In the above expression, $\boldsymbol{D}$ is a diagonal matrix with elements independently and randomly drawn from $\{-1, 1\}$. The matrix $\boldsymbol{H}$ represents a discrete Walsh–Hadamard transform, a generalization of a discrete Fourier transform over a modulo-2 number field. Both matrices are orthogonal and their action on $\boldsymbol{Q}$ does not affect the mutual distance structure; they simply perform rotations and reflections on the data sequence.

The action of $\boldsymbol{H}$ on a snapshot transforms the signal vector from physical to spectral space. Based on the uncertainty relation of harmonic analysis, a signal and its spectrum cannot both be concentrated: global signals have a narrow spectral content, while compact signals have broad spectral support. As a consequence, the $\boldsymbol{H}$ matrix maps narrow signals to dense vectors and produces sparse vectors from dense (e.g. single-scale) signals. In view of the fact that $\boldsymbol{P}$ will extract a very small number of components from the $\boldsymbol{H}$-transformed signals, we are interested in producing dense vectors – even for single-scale signals – to avoid a zero output vector from the FJLT. To circumvent sparse output vectors altogether, the random matrix $\boldsymbol{D}$ is added. More mathematically, it can be shown (Ailon & Chazelle 2010) that the largest $\boldsymbol{HD}$-transformed component of a unit-norm flow field $\boldsymbol{q}$, i.e. $\|\boldsymbol{HDq}\|_\infty$, is bounded by $O(\sqrt{\log N}/\sqrt{d})$, with a failure probability of less than $1/20$. This procedure is referred to as a spectral densification, as it always yields a dense vector, regardless of whether the input signal is single-scale or broadband in nature.

Once a proper densification has been established by $\boldsymbol{HD}$, a random projection via $\boldsymbol{P}$ can then accomplish the model reduction in a reliable and efficient manner. The final matrix $\boldsymbol{P}$ is given by

$$\boldsymbol{P} = \mathtt{Bernoulli}(q)\mathcal{N}(0, q^{-1}) = \begin{cases} \mathcal{N}(0, q^{-1}) & \text{with prob. } q \\ 0 & \text{with prob. } 1-q, \end{cases} \tag{2.4}$$

with $\mathtt{Bernoulli}(q)$ as a Bernoulli distribution, taking on the value of one with probability q and zero with probability $1 - q$. The normal distribution with zero mean and variance σ is denoted by $\mathcal{N}(0, \sigma)$. The parameter q measures the approximate sparsity

of $\boldsymbol{P}$. Detailed analysis requires q to be

$$q = \min\left(\mathcal{O}\left(\frac{\log^2 N}{d}\right), 1\right). \tag{2.5}$$

The compression that can be achieved with this parameter set-up can be estimated as (Dasgupta & Gupta 1999, 2003)

$$k \geq 4\left(\frac{\varepsilon^2}{2} - \frac{\varepsilon^3}{3}\right)^{-1} \log N, \tag{2.6}$$

which, for small values of the distortion factor ε, reduces to $k \sim \varepsilon^{-2} \log N$. Commonly, this reduction to a latent-space dimension of k is driven by a bound on the failure probability δ which, in turn, produces the above estimate. It is noteworthy that the reduced latent-space dimension k is independent of the input dimension d, but rather only scales with the enforced distortion factor and the imposed linkage parameter.

Most algorithms used in the processing of data sequences scale geometrically with the input dimension d and the snapshot number n; for example, the singular value decomposition (SVD) of a tall-and-skinny data matrix has a $O(dn^2)$ operation count. This scaling quickly brings these algorithms to the breaking point of applicability, a phenomenon often referred to as the curse of dimensionality. The above compression counteracts this tendency and ultimately yields a more favourable scaling which is amenable to large-scale applications. For this reason, data-compression and hashing algorithms (like the FJLT) have entered iterative linear-algebra algorithms as data preconditioners (Nelson & Nguyên 2013). In these applications, the dynamics preservation translates to the maintenance of relational connections between subsequent members of a Krylov sequence, and maintenance of spectral information. In this study, we apply the FJLT technique to dynamic flow-field sequences.

The implementation of the FJLT procedure does not involve the explicit formation of the three matrix components $\boldsymbol{P}$, $\boldsymbol{H}$, $\boldsymbol{D}$, nor does the entire data sequence $\boldsymbol{Q}$ have to be stored. Rather, (1) the multiplication of a snapshot by the diagonal matrix $\boldsymbol{D}$ amounts to randomly flipping the signs on the components of the input vector, (2) the multiplication by $\boldsymbol{H}$, the only dense matrix, can be accomplished in $O(d \log d)$ operations, similar to a fast Fourier transform, and (3) the multiplication by $\boldsymbol{P}$ is very efficient due to its substantial sparsity; it practically amounts to a small random sampling/multiplication of the output from $\boldsymbol{HDq}_i$. The FJLT can even be applied in a streaming fashion to the data sequence, which also permits the application of an incremental version of the SVD (see e.g. Brand 2002) for the subsequent POD or DMD analysis.

3. Problem set-up

To demonstrate the functionality and utility of data compression and model reduction by FJLTs, we make use of a model flow through a turbomachinery compressor blade row. This flow has previously been analysed in Glazkov (2021) and Glazkov *et al.* (2023*a*,*b*) as to its global stability and receptivity characteristics. The flow configuration consists of a single, two-dimensional blade passage, formed by controlled-diffusion airfoils, with periodic boundary conditions forming a linear blade row. The purely two-dimensional flow field is simulated for a chord-based Reynolds number of $Re = 100\,000$, with an exit Mach number of $M = 0.3$, and an inflow angle of 37.5°, which, downstream of the blade, is turned into a streamwise flow.

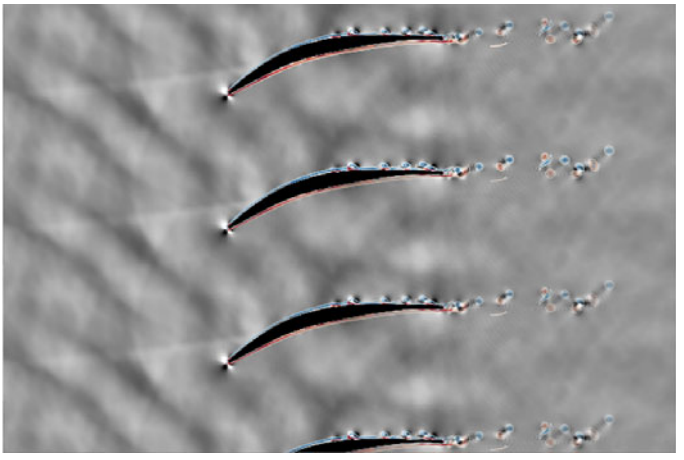

Figure 1. The dilatation field of a nonlinear flow snapshot for this flow case, upon which are superimposed contours of vorticity.

In this configuration the main features of the flow field can be summarized as comprising a laminar separation bubble on the suction side of the airfoils which periodically sheds vortices that subsequently interact with the trailing edge. These interactions of scattering and shedding processes at the trailing edge generate upstream-propagating acoustic waves that impinge upon the unstable boundary layers on the suction and pressure surfaces of the airfoil, as shown in the nonlinear flow snapshot in figure 1. Through a detailed wavemaker analysis in Glazkov (2021) and Glazkov *et al.* (2023*a*), it was shown that these interactions lead to feedback loops that sustain and drive the instabilities, and an ensemble of 14 dominant modes were identified and analysed through a direct-adjoint mean-flow global stability analysis. With such richness of dynamics, this case provides us with a sound foundation upon which to evaluate the performance of the FJLT when applied to modal flow analysis techniques explored in this paper.

The data set considered here consists of time-resolved snapshots of the four flow variables p, s, u and v (pressure, entropy, streamwise and spanwise velocity), extracted from the simulation once the system reaches a limit cycle. A total of $n = 2001$ snapshots are recorded over the time period $t = 0 \rightarrow 40$, and each flattened snapshot has $d = 1\,804\,400$ degrees of freedom. As is standard for POD and DMD, a $d \times n$ data matrix $\boldsymbol{Q}$ is then formed by consecutively stacking each snapshot into the columns of the matrix.

For a POD analysis, $\boldsymbol{Q}$ is decomposed using the SVD according to

$$\boldsymbol{Q} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{H}, \tag{3.1}$$

where we emphasize that $\boldsymbol{U}$ encodes spatial information in the form of POD modes, while $\boldsymbol{\Sigma}$ and $\boldsymbol{V}^{H}$ are purely temporal, encoding the singular values (spectrum) and evolution behaviour, respectively. Since the FJLT only acts on the spatial dimension of $\boldsymbol{Q}$, it follows that, in the latent (reduced) space, we have

$$\boldsymbol{B} = \boldsymbol{F}\boldsymbol{\Sigma}_{c}\boldsymbol{V}_{c}^{H}, \tag{3.2}$$

where

$$\boldsymbol{B} = \mathtt{FJLT}(\boldsymbol{Q}) \equiv \boldsymbol{PHDQ}, \tag{3.3}$$

and $\boldsymbol{B}$ is a $k \times n$ matrix with $k \ll d$, while $\boldsymbol{F}$ describes the POD modes in latent space. Although the inversion back to $\boldsymbol{U}$ is not possible directly from $\boldsymbol{F}$, the JL-guaranteed

Linkage N	Latent-space dimensions k	Compression (direct) d/k	Compression (delay-embedded) Nd/k	Failure probability δ
2	55 448	32.5	65.1	$\leq 1/3$
4	110 896	16.3	65.1	$\leq 1/18$
6	143 336	12.6	75.5	$\leq 1/45$
8	166 352	10.8	86.6	$\leq 1/84$
12	198 792	9.1	108.9	$\leq 1/198$
16	221 800	8.1	130.2	$\leq 1/360$
24	254 240	7.1	170.3	$\leq 1/828$
32	277 256	6.5	208.3	$\leq 1/1488$

Table 1. Compression ratios and failure probabilities for different FJLT embedding dimensions (linkage parameters).

closeness of $\boldsymbol{\Sigma}_c \boldsymbol{V}_c^H$ to $\boldsymbol{\Sigma} \boldsymbol{V}^H$ means that the inversion can be performed directly, though approximately, from (3.1) at a markedly reduced computational cost when compared with a direct SVD of $\boldsymbol{Q}$.

In a similar way, choosing DMD as our method of choice for a decomposition of $\boldsymbol{Q}$, we have

$$\boldsymbol{Q} = \boldsymbol{\Psi} \boldsymbol{A} \boldsymbol{V}_{\mathtt{and}}, \tag{3.4}$$

with $\boldsymbol{\Psi}$ signifying the spatial dynamic modes, $\boldsymbol{A}$ containing the amplitudes and $\boldsymbol{V}_{\mathtt{and}}$ carrying the temporal dynamics (in pure frequencies). Again, the FJLT approximately preserves the temporal dynamics contained in $\boldsymbol{A}\boldsymbol{V}_{\mathtt{and}}$. For the recovery of the spatial dynamic modes, the inversion of the Vandermonde matrix $\boldsymbol{V}_{\mathtt{and}}$, however, proves more challenging. Instead, in this paper, the physical modes $\boldsymbol{\Psi}$ are recovered using a discrete Fourier transform evaluated at the identified DMD frequencies and scaled by the DMD amplitudes. All spectral information about the retained dominant modes has been obtained using the sparsity-promoting variant of DMD (Jovanović, Schmid & Nichols 2014) with a penalization parameter selected to yield between 20 and 30 prevailing modes.

A related approach has been taken by Brunton *et al.* (2015) who used compressed-sensing techniques to preprocess the data stream that was subsequently analysed by DMD. In their study, Fourier sparsity is assumed and the temporal dynamics or spectral properties between the physical and latent space are not necessarily preserved.

4. Results

In all presented cases, the error tolerance has been set to a 1 % maximal distortion of the intrinsic dynamics, with the embedding dimension k determined by the number of snapshots N (Larson & Nelson 2017) over which this error tolerance is to be preserved (following the bounds given in Dasgupta & Gupta (1999, 2003)). In table 1, the values of N, considered in this study, are illustrated alongside the corresponding embedding dimension k, the compression ratio relative to a single snapshot, d/k, and the compression ratio relative to a column of the equivalently Hankelized system, Nd/k. We also list the failure probabilities as a function of the linkage parameter N. We note that increasing the linkage number N, and therefore k, increases the dimension of the column vectors in latent space, but relative to the Hankelized system, formed by a time-delay embedding over the same N state vectors, the compression ratio steadily increases, progressively favouring the FJLT compression technique over the equivalent Ruelle–Takens embedding.

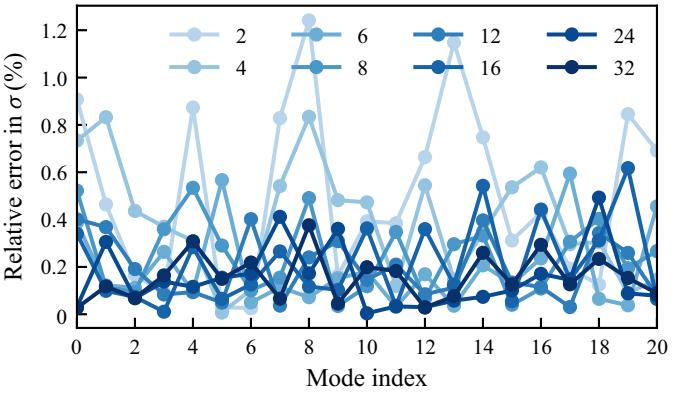

Figure 2. Relative errors in the singular values σ for a range of linkage parameter N. The first 20 POD modes have been computed, and a maximum distortion rate of 1 % has been chosen.

4.1. *Proper orthogonal decomposition (SVD/POD)*

To demonstrate the POD mode preservation, $n = 500$ temporal snapshots of the flow are considered in the $\boldsymbol{Q}$ matrix. From here on, the term 'direct' refers to quantities obtained directly from $\boldsymbol{Q}$, while the label 'FJLT' refers to those generated from $\boldsymbol{B}$, as defined in (3.3). Furthermore, we apply an economy-type SVD on the entire original or FJLT-reduced snapshot sequences. This choice isolates the influence of the compression algorithm on the full analysis. Alternative iterative SVD schemes, applied to each sequence, are left for a future effort.

Figure 2 displays the relative errors of the singular values, corresponding to the principal 20 POD modes for a range of linkage parameters N. For the lowest linkage of $N = 2$ we observe that the relative error occasionally exceeds the user-defined distortion error of 1 %. This behaviour is particularly prevalent for the higher modal indices, beyond the first four singular values. However, even with this crude approximation and sizeable compression, the relative error does not exceed a value of about 1.3 %. As we increase the linkage parameter N, the relative-error curve decreases markedly and remains (within our numerical experiments) below the threshold of 1 %. While there is no observable monotonic convergence towards lower errors, an overall tendency to smaller distortions, across the computed 20 POD modes, is clearly discernible.

For our second test, we concentrate on the temporal dynamics contained in the rows of the matrix $\boldsymbol{V}^H$. For a proper comparison, the dynamics of the direct case and the FJLT case have been Fourier-transformed and the difference between the respective two power spectra has been evaluated. Figure 3 depicts the ℓ_1 error of the difference vector in frequency space. Again, we observe a trend towards smaller relative errors as the linkage parameter N increases. More important, however, is the observation that the first four rows of $\boldsymbol{V}^H$ are well represented by the FJLT-compressed matrix $\boldsymbol{B}$ – even for very low values of N. Higher modes (beyond a modal index of four) require a higher linkage parameter N to achieve comparable relative ℓ_1 errors.

As a final test, we determine the six most dominant POD structures $\boldsymbol{U}_{1,\ldots,6}$ associated with the identified principal singular values. We choose a linkage number of $N = 6$ which has demonstrated in our previous test a sufficiently small relative error in the singular values as well as the temporal dynamics. As mentioned above, the spatial modes $\boldsymbol{U}$ cannot be directly recovered from the latent equivalents $\boldsymbol{F}$. However, since $\boldsymbol{\Sigma}_c \boldsymbol{V}_c^H$ is minimally distorted relative to the base case, we use the relation $\boldsymbol{U} = \boldsymbol{Q}\boldsymbol{V}_c\boldsymbol{\Sigma}_c^{-1}$, where both $\boldsymbol{V}_c$

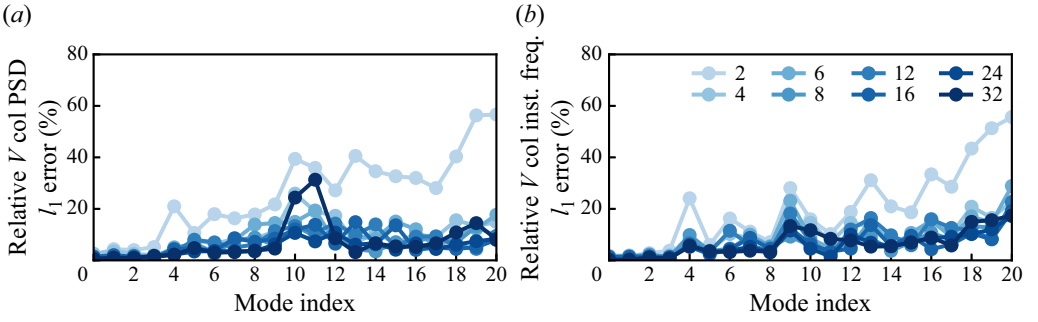

Figure 3. Relative ℓ_1 errors in the frequency domain of the column vectors of $\boldsymbol{V}$ obtained from the decomposition of $\boldsymbol{B}$ for various values of N. A maximum distortion rate of 1 % has been chosen.

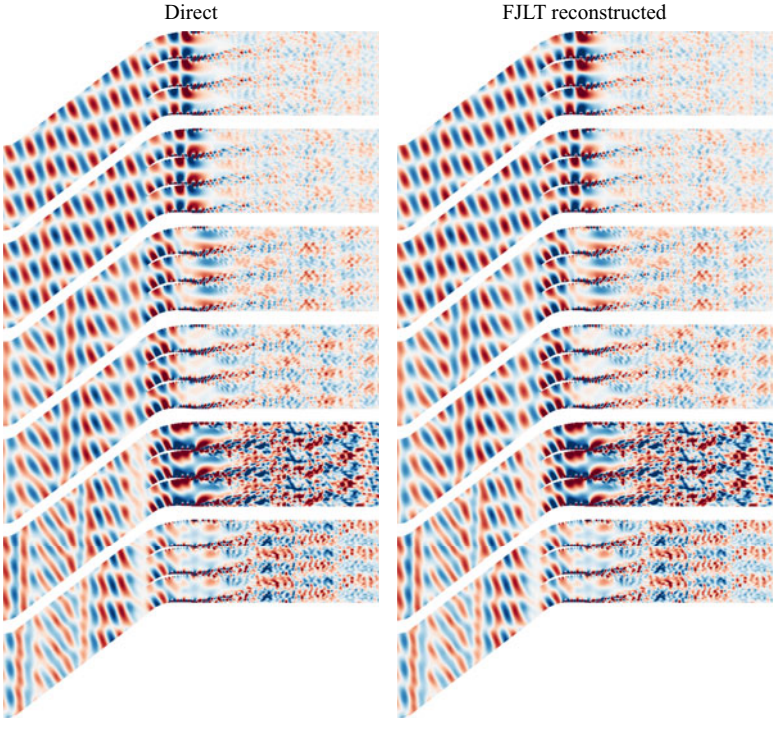

Figure 4. Proper orthogonal modes, corresponding to the six largest singular values, obtained from the direct POD (left column) and reconstructed from the FJLT-compressed snapshot sequence (right column). A linkage parameter of $N = 6$ and a maximum distortion factor $\varepsilon = 0.01$ have been chosen.

and $\boldsymbol{\Sigma}_c$ are from the FJLT case, to approximate the original spatial POD modes. The first six POD modes are displayed in figure 4, showing the baseline modes in the left column and the corresponding FJLT-based, reconstructed modes in the right column. We confirm a remarkable semblance between the two columns, which manifests itself even in the smaller details (see e.g. the complex wake structure of the fifth mode). We notice a distortion neither in amplitude nor in phase. The FLJT-based analysis, however, has been accomplished at a highly reduced computational cost (see the data in table 2).

Linkage N	Speed-up (POD) S_1	Speed-up (DMD) S_2	Speed-up (H-POD) $\approx NS_1$	Speed-up (H-DMD) $\approx NS_2$
2	33.18	26.26	66.38	52.53
4	19.84	15.36	79.37	61.44
6	15.98	11.50	95.89	69.01
8	14.26	10.55	114.04	84.42
12	12.15	9.66	145.76	115.92
16	10.78	8.79	172.40	140.59
24	9.38	7.04	225.07	168.92
32	8.73	6.55	279.36	209.44

Table 2. Timing comparison of direct and FJLT-driven analysis.

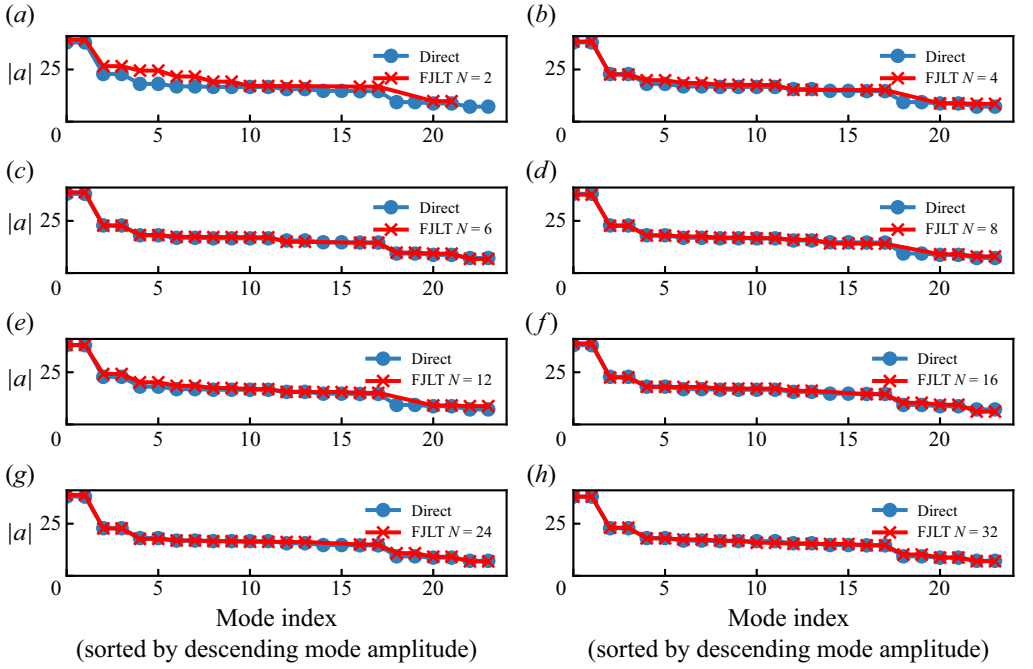

Figure 5. The spDMD amplitudes for (*a–h*) varying N.

4.2. *Sparsity-promoting DMD (spDMD)*

Using the same data sequence and performing a flow analysis based on the spDMD, coupled to our FJLT compression, also shows encouraging results at a highly reduced computational effort. The amplitudes, identified by the sparsity-promoting optimization, are displayed for the dominant 24 dynamic modes in figure 5 for the same range of parameter values for N. Excellent representation for the highest pair of structures can be observed, even for the minimal linkage of $N = 2$. Lower-amplitude, subdominant modes are well represented as the linkage parameter N is increased, and even modest values of N accurately capture the amplitudes of the majority of the top 24 dynamic modes.

The same tendency can be observed in spectral space. Figure 6 displays the eigenvalues of the recovered Koopman operator in the complex plane, again for our range of linkage parameters N. A part of the unit disk is displayed for each case, showing the direct

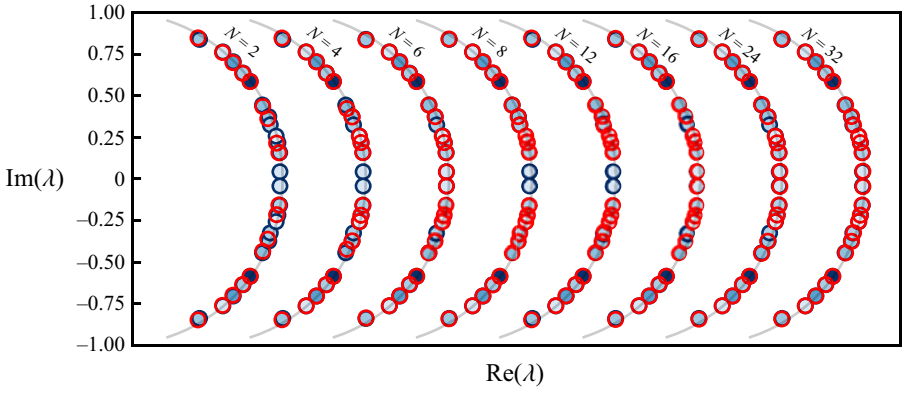

Figure 6. Direct and reconstructed DMD eigenvalues. Blue filled circles indicate the direct eigenvalues, with darker circles representing higher amplitudes, and red circles indicating the corresponding FJLT-derived eigenvalues for varying N.

eigenvalues in blue and the FJLT-based eigenvalues in red. In addition, the colouring encodes the associated amplitude of each eigenvalue, with darker shades corresponding to larger amplitudes. We see that the eigenvalues close to the dominant structure (darkest symbol) are well represented by the FJLT equivalent, even for a coarse linkage of $N = 2$. The lower-frequency structures, near the spectral point $\lambda = (1, 0)$, require a higher linkage parameter N, before they can be extracted reliably from the FJLT-compressed data set, but are nonetheless well represented when compared with their direct equivalents.

A more quantitative analysis of the eigenvalue errors is presented in figure 7. We stress that the user-specified distortion of 1 % in phase-space distance does not translate in a straightforward manner into an estimate for the DMD eigenvalue deviation. Nevertheless, we notice a pleasing tendency of the eigenvalue error to reduce for increasing values of N. The eigenvalues associated with the largest two amplitudes (representing the dominant dynamic structures in the processed data sequence) show very low relative errors with values below 0.15 %, even for the crudest of approximations ($N = 2$). Less dominant structures, i.e. higher modes, representing mostly lower-frequency structures are represented with a maximum eigenvalue error of about 1 % once the linkage parameter exceeds $N = 6$. Even at this point, the computational savings, compared with the direct case, are substantial.

The reconstruction of the spatial structures corresponding to the identified eigenvalues/frequencies is less straightforward than for the POD analysis. A technique similar to the POD case would require the inversion of a Vandermonde matrix and compound the FJLT approximation error with additional numerical errors stemming from the inversion. For this reason, we chose a different reconstruction method, using the fact that for statistically equilibrated data sequences the eigenvalues are located on the unit disk and represent single-frequency structures. We thus identify the FJLT-compressed frequency and inverse-Fourier-transform the data sequence for this frequency. The results are shown in figure 8, juxtaposing the six direct DMD modes with the largest amplitudes, the equivalent six modes from a direct reconstruction and the corresponding modes from an FJLT-based recovery. While the recovered modes (direct and FJLT-based) match closely, corroborating the efficacy of the compression algorithm to capture significant flow characteristics from the latent data sequence, there are noticeable differences, mostly in phase, between the direct and reconstructed coherent structures. Nevertheless, the key

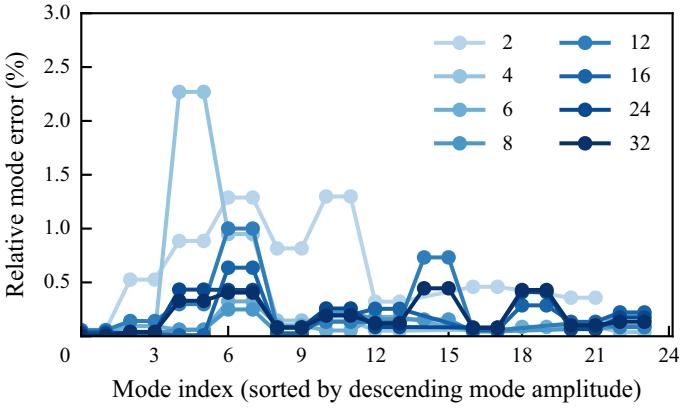

Figure 7. The relative errors in the FJLT-reconstructed modes given by $err = |\lambda_{dir} - \lambda_{FJLT}|/|\lambda_{dir}|$.

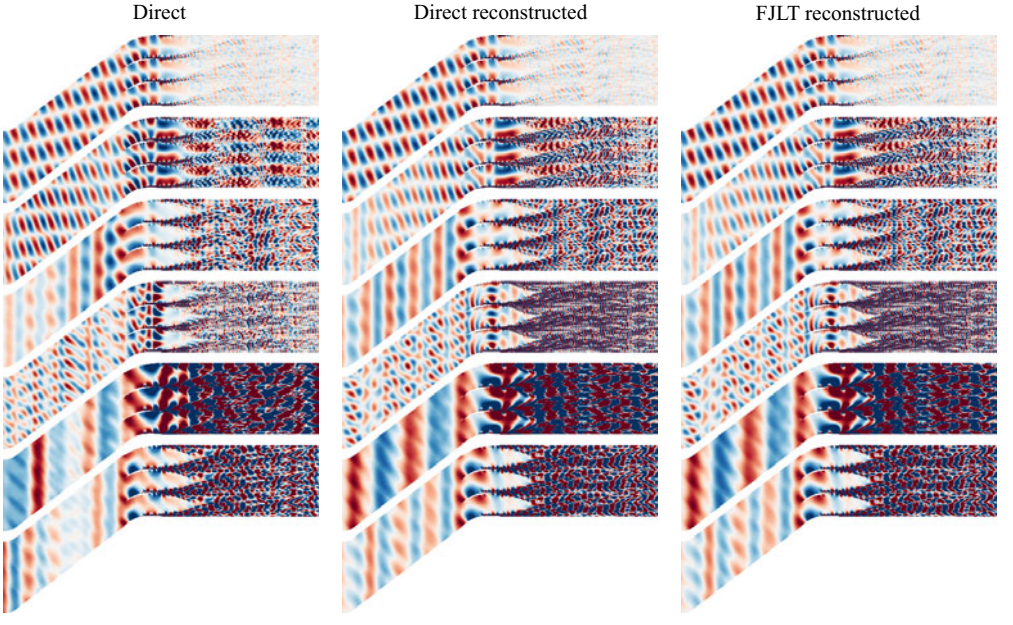

Figure 8. A DMD reconstruction from an FJLT sequence. Shown are the six most dominant dynamic modes, with the reference case on the left, and the reconstructed modal structures in the middle and on the right. Contours of pressure are used to visualize the flow fields.

features in each of the six modes are represented in a very satisfactory manner. It is important to mention that the differences can be attributed to the reconstruction, not the compression procedure, and should be taken as an encouragement to direct further attention to a better extraction method for the spatial structure from the latent dynamics.

4.3. *Computational cost, timings and failure rates*

The logarithmic scaling of the latent-space dimensionality translates to an equivalent logarithmic scaling in execution time, since the SVD of a rectangular matrix scales linearly with the larger of the two dimensions. With the execution time of the base case established

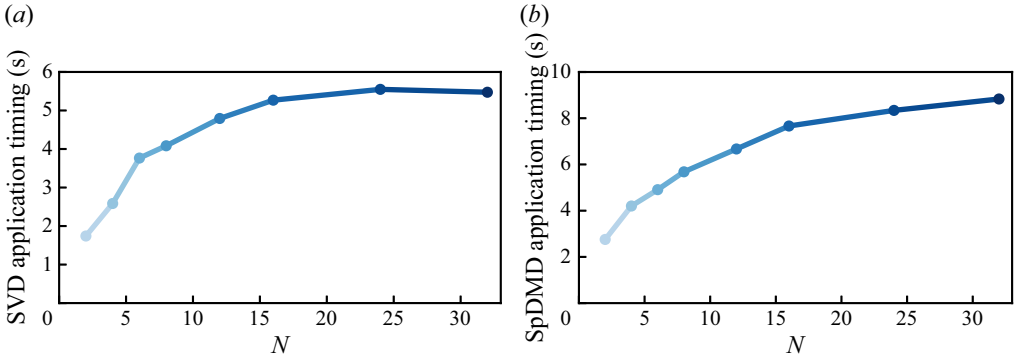

Figure 9. Timing of the FJLT-based POD (*a*) and spDMD (*b*) versus the linkage number N. We note that the embedding dimension k scales logarithmically with the linkage number N; consequently, the horizontal axis can be interpreted as the embedding dimension.

and fixed, we determine the speed-up factor of the FJLT-based modal decomposition as the linkage parameter N is varied from $N = 2$ to $N = 32$. For cases of both POD and spDMD analyses, we corroborate an inverse logarithmic decrease in the speed-up factor (see table 2). When multiplying by the embedding dimension for the processing of a delay-embedded (Hankelized) data sequence, the speed-up factors become even more impressive, asymptotically displaying an $N/\log N$ tendency (see the columns in table 2 labelled H-POD and H-DMD). Figure 9 depicts the raw execution speeds (executed on a 32 GB-RAM, Apple M1 Max laptop) versus the linkage parameter graphically, and with the colour code of earlier figures. Both cases (POD and spDMD analyses) are presented, and a clear logarithmic scaling is confirmed.

In a final analysis, we address the failure rate of the FJLT bound. As mentioned earlier, the bounds on the distortion (2.1), on which our accelerated modal analysis rests, have to be understood stochastically. In other words, the bounds are satisfied up to a failure rate δ, when the actual distortion exceeds the specified values of $(1 \pm \varepsilon)$. The rate at which this outlier occurs gives a measure of reliability of our compression algorithm. For this reason, we present a statistical experiment where we compute the distortion in the pairwise distance between randomly sampled snapshots from our data sequence, i.e. we compute

$$\varepsilon_{true} = \left| \frac{\|\boldsymbol{\Phi}\,\Delta\boldsymbol{q}\|_2}{\|\Delta\boldsymbol{q}\|_2} - 1 \right|, \tag{4.1}$$

with $\Delta\boldsymbol{q} = \boldsymbol{q}_i - \boldsymbol{q}_j$, $i, j \in [0, n]$ and process the data into a histogram. Again, we select a desired value of $\varepsilon = 1\,\%$ and perform the experiment for a linkage parameter of $N = 2, 4, 6$. We sampled our data sequence 60 000 times (without repetition) to obtain properly converged statistics. The results are shown in figure 10. For $N = 2$ we observe a failure rate of 22.37 %; in other words, nearly a quarter of our samples exceeded the desired distortion of $\varepsilon = 1\,\%$. This value confirms the upper bound of $\delta = 1/3$. The distribution in figure 10, however, also shows that distortions larger than 2 % occur rarely, even for this minimal value of $N = 2$. Increasing N improves the failure rate dramatically. Already for $N = 4$, only 1.48 % of our samples exceed the user-specified 1 % distortion value, and at $N = 6$, the failure rate has reduced to 0.0156 %. This experiment also shows that the theoretical bounds on the failure rate, given in the rightmost column of table 1, are rather conservative, as they predict a maximum failure rate of $1/18 = 5.55\,\%$ for $N = 4$

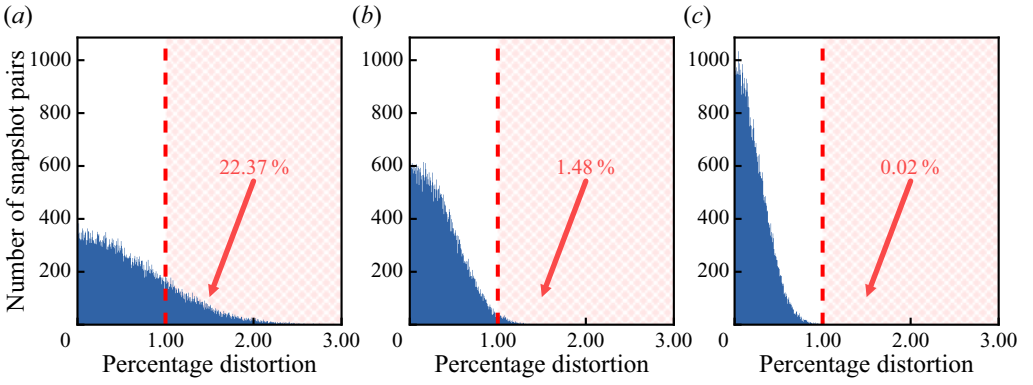

Figure 10. Error (failure rate) histograms from a statistical analysis of the FJLT-based POD for a linkage number of (*a*) $N = 2$, (*b*) $N = 4$ and (*c*) $N = 6$. The respective upper bounds for the failure probability are 33.33 % for $N = 2$, 5.55 % for $N = 4$ and 2.22 % for $N = 6$.

and $1/45 = 2.22\,\%$ for $N = 6$. Our statistical analysis produced substantially lower failure rates in both cases.

5. Summary and outlook

High spatial resolution, and hence many degrees of freedom, is a computational necessity for capturing all relevant spatio-temporal scales and physical processes when simulating complex fluid flows, but acts as an impediment to algorithmic efficacy when the resulting flow fields are subsequently analysed. Common efforts in reducing flow mechanisms to their relevant subprocesses via spectral or modal decompositions yield distortions or misrepresentations of the temporal dynamics and result in a latent-space dynamics that is difficult to interpret or relate to full-scale observations.

In this study, we propose a spatial reduction scheme that achieves up to two orders of magnitude compression ratios while maintaining the temporal dynamics within a predetermined distortion tolerance. This nearly isometric (distance-preserving) reduction takes advantage of the JL lemma and its fast implementation via the FJLT. Randomized sketching techniques of preprocessed flow fields provide a low-dimensional embedding and produce substantially reduced latent-space sequences that are further subjected to modal analyses at diminished computational costs.

Applications to compressible flow through a linear blade cascade demonstrated the efficacy of the algorithm, where the dominant POD and DMD modes have been accurately, but far more efficiently, captured. Singular values, frequencies, amplitudes and modal shapes could be extracted from reduced data sets without significant degradation in quality despite a data compression factor of up to $O(10^2)$.

The algorithm is governed by a linkage parameter that enforces the distortion bounds across a set of snapshots. For the special case of successive linkage, this parameter can be interpreted as a time horizon over which temporal coherence is enforced during the compression. Our proposed algorithm can thus be considered as an alternative to time-delay (Ruelle–Takens) embedding. However, while time-delay embedding substantially inflates the spatial dimensionality of an already high-dimensional state vector, we venture in the opposite direction by reducing the spatial dimensionality while preserving temporal structures over the same time horizon. Our proposed algorithm

exploits the compressibility (in a data-information sense) of the flow field, rather than introducing excessive and unnecessary redundancy by Hankelizing the data matrix.

While the FJLT-based flow analysis produced encouraging results, there still remain opportunities for improvements and extensions. Parallelization efforts come to mind to further increase the efficiency of the data compression. While the Walsh–Hadamard transform can be parallelized in itself, a more direct gain in execution speed may arise from a partitioning of the snapshot sequence and the parallel processing of each subset. Throughout this paper we have made use of the original approach of Ailon & Chazelle (2009), Ailon & Chazelle (2010) and Ailon & Liberty (2013), but recent work on FJLTs (see e.g. Fandina, Høgsgaard & Larsen 2022) has sought to improve the computational complexity associated with the application of the transform. This would be of particular relevance when, for example, applying the transform in the contexts of long time series where the transform must be applied many times. The measure-densifying effect of the $\boldsymbol{HD}$ matrix (in our case, a Hadamard–Walsh transform of a sign-randomized signal) can be replaced by other options that accomplish the same densification. In particular, in-place transforms could be applied to alleviate the memory footprint of the compression. Most importantly, recovery techniques to reconstruct modal structures from their latent-space analogues are a target of future research, with an eye on machine-learning techniques, such as superresolution, to tackle this challenge. However, even in its original form presented here, the FJLT-based modal flow analysis represents an efficient and appealing way of gaining physical insight into high-dimensional, large-scale and complex flow processes, without the accompanying computational effort.

Supplementary material. Codes and data sets to reproduce the results in this paper are provided as supplementary material, available at https://github.com/antongla/fjlt-modal-compression.

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

Author ORCIDs.
Anton Glazkov https://orcid.org/0000-0003-0137-199X;
Peter J. Schmid https://orcid.org/0000-0002-2257-8490.

References

Achlioptas, D. 2003 Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. Syst. Sci.* **44** (4), 671–687.

Adrian, R. & Moin, P. 1988 Stochastic estimation of organized turbulent structure: homogeneous shear flow. *J. Fluid Mech.* **190**, 531–559.

Ailon, N. & Chazelle, B. 2009 The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. *SIAM J. Comput.* **39** (1), 302–322.

Ailon, N. & Chazelle, B. 2010 Faster dimension reduction. *Commun. ACM* **53** (2), 97–104.

Ailon, N. & Liberty, E. 2013 An almost optimal unrestricted fast Johnson-Lindenstrauss transform. *ACM Trans. Algorithms* **9** (3), 1–12.

Arbabi, H. & Mezić, I. 2017 Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the Koopman operator. *SIAM J. Appl. Dyn. Syst.* **16** (4), 2096–2126.

Axås, J. & Haller, G. 2023 Model reduction for nonlinearizable dynamics via delay-embedded spectral submanifolds. *Nonlinear Dyn.* **111** (24), 22079–22099.

Berkooz, G., Holmes, P. & Lumley, J. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25**, 539–575.

Brand, M. 2002 Incremental singular-value decomposition of uncertain data with missing values. In *Proceedings of the 7th European Conference on Computer Vision*, pp. 707–720. Springer.

Brunton, S.L., Brunton, B.W., Proctor, J., Kaiser, E. & Kutz, J.N. 2017 Chaos as an intermittently forced linear system. *Nat. Commun.* **8** (19), 1–9.

Brunton, S.L., Proctor, J.L., Tu, J.H. & Kutz, J.N. 2015 Compressed sensing and dynamic mode decomposition. *J. Comput. Dyn.* **2** (2), 165–191.

BUDISIĆ, M., MOHR, R. & MEZIĆ, I. 2012 Applied Koopmanism. *Chaos* **22** (4), 047510.
CANDÈS, E.J. & TAO, T. 2006 Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inf. Theory* **52** (12), 5406–5425.
DASGUPTA, S. & GUPTA, A. 1999 An elementary proof of the Johnson-Lindenstrauss lemma. *Tech. Rep.* 99-006. UC Berkeley.
DASGUPTA, S. & GUPTA, A. 2003 An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Struct. Algorithms* **22** (1), 60–65.
FANDINA, O.N., HØGSGAARD, M.M. & LARSEN, K.G. 2022 The fast Johnson-Lindenstrauss transform is even faster. arXiv:2204.01800.
GLAZKOV, A. 2021 A direct-adjoint framework for stability and sensitivity analyses of turbomachinery aeroacoustics. PhD thesis, University of Oxford.
GLAZKOV, A., FOSAS DE PANDO, M., SCHMID, P.J. & HE, L. 2023*a* Global stability analysis of an idealized compressor blade row. I. Single-blade passage analysis. *Phys. Rev. Fluids* **8**, 103903.
GLAZKOV, A., FOSAS DE PANDO, M., SCHMID, P.J. & HE, L. 2023*b* Global stability analysis of an idealized compressor blade row. II. Multiple-blade interactions. *Phys. Rev. Fluids* **8**, 103904.
GOLYANDINA, N., NEKRUTIN, V. & ZHIGLJAVSKY, A. 2001 *Analysis of Time Series Structure: SSA and Related Techniques*. Chapman & Hall/CRC.
HALKO, N., MARTINSSON, P.G. & TROPP, J.A. 2011 Finding structure in randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* **53** (2), 217–288.
HOLMES, P., LUMLEY, J., BERKOOZ, G. & ROWLEY, C.W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press.
HYVÄRINEN, A., KARHUNEN, J. & OJA, E. 2001 *Independent Component Analysis*. John Wiley & Sons.
JOHNSON, W.B. & LINDENSTRAUSS, J. 1984 Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Maths* **26**, 189–206.
JOVANOVIĆ, M., SCHMID, P.J. & NICHOLS, J.W. 2014 Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **14**, 024103.
KAMB, M., KAISER, E., BRUNTON, S.L. & KUTZ, J.N. 2020 Time-delay observables for Koopman: theory and applications. *SIAM J. Appl. Dyn. Syst.* **19** (2), 886–917.
KUTZ, J.N., BRUNTON, S.L., BRUNTON, B.W. & PROCTOR, J.L. 2016 *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM Publishing.
LARSON, K.G. & NELSON, J. 2017 Optimality of the Johnson-Lindenstrauss lemma. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, pp. 633–638. IEEE.
LAWRENCE, N. 2012 A unifying probabilistic perspective for spectral dimensionality reduction: insights and new models. *J. Mach. Learn. Res.* **13**, 1609–1638.
LUMLEY, J. 1970 *Stochastic Tools in Turbulence*. Academic Press.
MARTINSSON, P.G. & TROPP, J.A. 2020 Randomized numerical linear algebra: foundations and algorithms. *Acta Numerica* **29**, 403–572.
MCKEON, B.J. 2017 The engine behind (wall) turbulence: perspectives on scale interactions. *J. Fluid Mech.* **817**, P1.
NELSON, J. & NGUYÊN, H.L. 2013 OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 117–126. IEEE.
PAN, S. & DURAISAMY, K. 2020 On the structure of time-delay embedding in linear models of non-linear dynamical systems. *Chaos* **30**, 073135.
RIBEIRO, J.H.M., YEH, C.A. & TAIRA, K. 2020 Randomized resolvent analysis. *Phys. Rev. Fluids* **5** (3), 033902.
RIGAS, G., SIPP, D. & COLONIUS, T. 2021 Nonlinear input/output analysis: application to boundary layer transition. *J. Fluid Mech.* **911**, A15.
ROWLEY, C.W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D.S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
SCHMID, P.J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
SCHMID, P.J. 2022 Dynamic mode decomposition and its variants. *Annu. Rev. Fluid Mech.* **54**, 225–254.
SCHMIDT, O.T. & COLONIUS, T. 2020 Guide to spectral proper orthogonal decomposition. *AIAA J.* **58** (3), 1023–1033.
SCHMIDT, O.T. & SCHMID, P.J. 2019 A conditional space-time POD formalism for intermittent and rare events: example of acoustic bursts in turbulent jets. *J. Fluid Mech.* **867**, R2.
SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. *Q. J. Appl. Maths* **45**, 561–590.

TAIRA, K., BRUNTON, S.L., DAWSON, S.T.M., ROWLEY, C.W., COLONIUS, T., MCKEON, B.J., SCHMIDT, O.T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L.S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4041.

TAKENS, F. 1981 Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence: Proceedings of a Symposium Held at the University of Warwick 1979/80* (ed. D. Rand & L.-S. Young), pp. 366–381. Springer.

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.* **647**, 821–867.