

Spectral Proper Orthogonal Decomposition using Multitaper Estimates

Oliver T. Schmidt¹

*Department of Mechanical and Aerospace Engineering, University of California San Diego,
La Jolla, CA 92093, USA*

Abstract

The use of multitaper estimates for spectral proper orthogonal decomposition (SPOD) is explored. Multitaper and multitaper-Welch estimators that use discrete prolate spheroidal sequences (DPSS) as orthogonal data windows are compared to the standard SPOD algorithm that exclusively relies on weighted overlapped segment averaging, or Welch's method, to estimate the cross-spectral density matrix. Two sets of turbulent flow data, one experimental and the other numerical, are used to discuss the choice of resolution bandwidth and the bias-variance tradeoff. Multitaper-Welch estimators that combine both approaches by applying orthogonal tapers to overlapping segments allow for flexible control of resolution, variance, and bias. At additional computational cost but for the same data, Multitaper-Welch estimators provide lower variance estimates at fixed frequency resolution or higher frequency resolution at similar variance compared to the standard algorithm.

Keywords:

1. Introduction

SPOD is the specialization of the most general form of proper orthogonal decomposition (POD, [1]) to statistically stationary data and is best computed in the frequency domain. An early application of the method can be found in the

¹oschmidt@ucsd.edu

5 experimental work of Glauser et al. [2]. The method has recently gained pop-
 ularity as theoretical connections to dynamic mode decomposition (DMD,[3])
 and resolvent analysis of turbulent mean flows [4] were established by Towne
 et al. [5]. Another factor is its application to large numerical data that were
 specifically generated for this purpose [6]. SPOD has been applied to a broad
 10 range of canonical [7, 8, 9, 10, 11, 12, 13] and technical [14, 15, 16, 17] turbulent
 flows, weather and climate data [18], and found diverse applications in aeroa-
 coustics [19, 20], stochastic estimation [21], frequency-time analysis [22], flow
 field reconstruction [23, 22], and reduced-order modeling [24, 25].

The SPOD modes and their mode energies are, respectively, the eigenvectors
 15 and eigenvalues of the weighted cross-spectral density (CSD) matrix. The CSD
 matrix is estimated from an ensemble of realizations of the temporal discrete
 Fourier transform (DFT). In contrast to classical Fourier analysis (see, e.g., Pain
 et al. [26] for a recent application to large data), SPOD modes are, in an inner
 product norm of choice, optimal linear combinations of many Fourier modes.
 20 The standard approach to estimate the CSD matrix for SPOD is weighted over-
 lapped segment averaging (WOSA), or Welch’s method [27, 28]. A different
 approach that is routinely used for the estimation of power spectra from time
 signals is the multitaper (MT), or Thomson’s multitaper method [29]. The mul-
 titaper method corresponds to weighted averaging of independent spectral esti-
 25 mates that are obtained from multiple orthogonal data windows. Most notably,
 it has been demonstrated by Bronez [30] that the multitaper power spectrum
 estimator outperforms the Welch estimator in terms of leakage, variance, and
 resolution if two of these three quantities are fixed to determine the third. In
 the context of spectral analysis of turbulent flow data, a frequency–wavenumber
 30 analysis technique that uses multitaper estimates was proposed by Geoga et al.
 [31].

The remainder of the paper is organized as follows. The technical back-
 ground is discussed in §2 and includes a brief summary of the discrete SPOD
 problem (§2.1) and DPSS (§2.2), as well as a recapitulation of Welch’s method,
 35 and the introduction of multitaper estimates for SPOD (§2.3). The two example

data sets are presented in §2.4. In §3, SPOD analyses conducted using the basic Welch and single-block multitaper estimates (§3.1) are compared and discussed in the light of parameter selection, frequency resolution, bandwidth, and the variance-bias tradeoff. Next, multitaper-Welch estimates that combine the advantages of both methods are computed for the two examples and discussed in §3.2. The additional computational cost associated with multitaper estimates is addressed in §4. Finally, implications and future applications of this work are discussed in §5.

2. Background and Methodology

2.1. Spectral proper orthogonal decomposition (SPOD)

The discrete SPOD of an ensemble of N_t snapshots of spatio-temporal multivariate data with zero mean,

$$\mathbf{q}_i = \mathbf{q}(t_i), \quad i = 1, \dots, N_t, \quad \mathbf{q} \in \mathbb{R}^{N_{\text{dof}} \times 1} \quad (1)$$

here written as a sequence of column vectors, is obtained from the eigendecompositions

$$\mathbf{C}_j \mathbf{W} \Phi_j = \Phi_j \Lambda_j, \quad j = 1, \dots, N_{\text{fft}}, \quad \mathbf{C}_j \in \mathbb{R}^{N_{\text{dof}} \times N_{\text{dof}}}, \quad (2)$$

of the cross-spectral density (CSD) matrices, \mathbf{C}_j . The rank M of the CSD matrices corresponds to the number of independent realizations of the Fourier transform that are used for their construction, and their size is determined by the number of degrees of freedom (number of variables times number of points in space), N_{dof} . The number of eigenvalue problems to be solved corresponds to the number of discrete frequencies, N_{fft} , of the discrete Fourier transform in the permissible frequency band between zero and the Nyquist frequency, $f_{\text{Nyq}} = \frac{f_s}{2}$, where $f_s = 1/\Delta t$ is the sampling frequency. The columns of $\Phi_j = [\phi_j^{(1)}, \phi_j^{(2)}, \dots, \phi_j^{(M)}]$, i.e., the eigenvectors of the weighted CSD matrices $\mathbf{C}_j \mathbf{W}$, are the SPOD modes and the diagonal entries of Λ_j , i.e., the corresponding eigenvalues $\lambda_f^{(1)} \geq \lambda_f^{(2)} \geq \dots \geq \lambda_f^{(M)}$, the mode energies. The weight matrix \mathbf{W}

accommodates spatial integration and, for multivariate data, variable-dependent weights. The discrete (sample) spatial CSD matrix is given as

$$\mathbf{C}_j = \frac{1}{M} \sum_{k=1}^M \hat{\mathbf{q}}_j^{(k)} \left(\hat{\mathbf{q}}_j^{(k)} \right)^* \quad (3a)$$

$$= \frac{1}{M} \hat{\mathbf{Q}}_j \hat{\mathbf{Q}}_j^*, \quad (3b)$$

where $\hat{\mathbf{q}}_j^{(k)}$ is the k -th out of M realizations of the discrete Fourier transform at the j -th frequency. The matrix representation, equation (3b), uses the data matrix $\hat{\mathbf{Q}}_j = [\hat{\mathbf{q}}_j^{(1)}, \hat{\mathbf{q}}_j^{(2)}, \dots, \hat{\mathbf{q}}_j^{(M)}]$. The denominator is either M or $M - 1$, depending on whether the true mean, if available, or the sample mean is subtracted from the data, respectively. In the common case where $M < N_{\text{dof}}$, the smaller size $M \times M$ eigenvalue problem,

$$\mathbf{C}'_j \Psi_j = \Psi_j \Lambda'_j, \quad \Phi_j = \hat{\mathbf{Q}}_j \Psi_j, \quad j = 1, \dots, N_{\text{fft}}, \quad (4a)$$

$$\text{where } \mathbf{C}'_j = \frac{1}{M} \hat{\mathbf{Q}}_j^* \mathbf{W} \hat{\mathbf{Q}}_j, \quad \mathbf{C}'_j \in \mathbb{R}^{M \times M}, \quad (4b)$$

is the weighted temporal CSD matrix and $\Lambda'_j \in \mathbb{R}^{M \times M}$ the reduced matrix of non-zero eigenvalues, can be solved for the SPOD expansion coefficients, Ψ_j , instead of using equation (2) to obtain the SPOD modes directly. The desirable mathematical properties of the SPOD, its continuous formulation, and algorithmic implementations are discussed in detail elsewhere [5, 28].

2.2. Discrete Prolate Spheroidal Sequences (DPSS)




Figure 1: Leading seven DPSS for $N = 4096$ and $b_{\text{win}} = 4$ compared to a standard Hamming window. The DPSS are normalized by the maximum value of \mathbf{s}_1 for clarity.

Discrete prolate spheroidal sequences (DPSS), or Slepian sequences, are the discrete-time equivalent of prolate spheroidal wave functions [32, 33]. They are the optimal solutions to the spectral concentration problem that seeks orthogonal time sequences whose discrete Fourier transform is maximally localized in a frequency band $[-\delta f, \delta f]$, defined by the resolution half-bandwidth, δf . Given the window length, N , and resolution half-bandwidth, δf , a common way to find the DPSS is the eigendecomposition

$$\mathbf{H}\mathbf{s} = \gamma\mathbf{s}, \quad (5)$$

where the entries of the Hermitian matrix $\mathbf{H} = \mathbf{H}(\delta f, N)$ are given by

$$h_{ij} = \frac{\sin(2\pi\delta f(j-i))}{\pi(j-i)}, \quad i, j = 0, \dots, N-1. \quad (6)$$

The eigenvalues $1 > \gamma_0 > \gamma_1 > \dots > \gamma_{N-1} > 0$ correspond to the spectral concentrations of energy of the respective tapers, $\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_{N-1}$, in the frequency band $[-\delta f, \delta f]$. Since \mathbf{H} is Hermitian, the DPSS are mutually orthogonal. It can be shown that only $2\delta f N$ (Shannon number) of the spectral concentrations are close to unity. It is hence customary to use only the first $N_{\text{win}} = 2\delta f N - 1$ Slepian tapers as windows. In terms of the time-halfbandwidth product, $b_{\text{win}} = \delta f N$, this statement becomes

$$N_{\text{win}} = \lfloor 2b_{\text{win}} \rfloor - 1, \quad (7)$$

and guarantees that the Fourier transforms of all windows are well-concentrated in $[-\delta f, \delta f]$ to minimize spectral leakage to neighboring frequencies. The leading 55 seven DPSS for $N = 4096$ and $b_{\text{win}} = 4$ are compared to a standard Hamming window in figure 1.

2.3. Welch and multitaper estimates of the SPOD

The N independent realizations of the discrete Fourier transform that form the columns of the sample CSD matrix in equation (3a) are obtained in one of two ways. The first option is to obtain well-separated sequences, or blocks, $\mathbf{q}_j^{(k)}$ with $j = 1, \dots, N_{\text{fft}}$ and $k = 1, \dots, N_{\text{blk}}$, one at a time. Denote by N_{blk} the total number of blocks, each consisting of N_{fft} consecutive snapshots, equally

spaced in time by Δt . This scenario occurs when data acquisition at a high sampling rate is the bottleneck, as is the case in many experimental settings. The second possibility is that the data is acquired as a single, long time series of N_t snapshots, $\mathbf{q}_j = \mathbf{q}(t_j)$, with $j = 1, \dots, N_t$. This is often the case in numerical settings, where the overall runtime is the bottleneck. In this case, the long time series is segmented into N_{blk} blocks, $\mathbf{q}_j^{(k)} = \mathbf{q}_{j+(k-1)(N_{\text{fft}}-N_{\text{ovlp}})+1}$ with $j = 1, \dots, N_{\text{fft}}$ and $k = 1, \dots, N_{\text{blk}}$, under the ergodicity hypothesis. Adjacent blocks are allowed to overlap by N_{ovlp} snapshots, which results in a total number of $N_{\text{blk}} = \left\lfloor \frac{N_t - N_{\text{ovlp}}}{N_{\text{fft}} - N_{\text{ovlp}}} \right\rfloor$ blocks. This segmentation is the idea behind the method of Welch [27], that was originally devised for the estimation of power spectra of time signals. In the same work, Welch established from theoretical considerations the best practice of using an overlap of 50%. We follow this best practice but note that the optimal overlap depends on the window function, see Heinzel et al. [34] for a comprehensive discussion. After segmenting the data, N_{blk} realizations of the Fourier transform are obtained from the weighted temporal discrete Fourier transform of each block,

$$\hat{\mathbf{q}}_j^{(k)} = \sum_{i=0}^{N_{\text{fft}}-1} w[i] \mathbf{q}_i^{(k)} e^{\frac{i2\pi}{N_{\text{fft}}} ij}, \quad k = 1, \dots, N_{\text{blk}}, \quad (8)$$

where $\hat{\mathbf{q}}_j^{(k)}$ is the k -th realization of the Fourier transform at the j -th discrete frequency. It is common practice to use a data window, \mathbf{w} , to prevent spectral leakage due to the non-periodicity of the samples. A popular choice is the symmetric Hamming window,

$$w[i] = 0.54 - 0.46 \cos\left(\frac{2\pi i}{N_{\text{fft}} - 1}\right), \quad i = 0, \dots, N_{\text{fft}} - 1. \quad (9)$$

Given the ensemble of realizations, equation (8), the CSD matrix is formed according to equation (3a), and its eigendecomposition, equation (2), concludes the standard SPOD algorithm.

Inspired by the method of Thomson [29], also originally intended for the estimation of power spectra, we propose to use a set of N_{win} orthogonal DPSS (see §2.2) as data windows for the SPOD. We restrict this set to the first $2b_{\text{win}} - 1$

sequences (equation (7)) to ensure that all data windows are well-concentrated in $[-\delta f, \delta f]$. Since the DPSS are orthogonal, the different realizations are independent. By combining segmenting and orthogonal windowing, a total number of $N_{\text{blk}}N_{\text{win}}$ realizations,

$$\hat{\mathbf{q}}_j^{(k,l)} = \sum_{i=0}^{N_{\text{fft}}-1} s_l[i] \mathbf{q}_i^{(k)} e^{\frac{i2\pi}{N_{\text{fft}}} ij}, \quad k = 1, \dots, N_{\text{blk}}, \quad l = 0, \dots, N_{\text{win}} - 1, \quad (10)$$

is obtained. Here, $s_l[i]$ is the i -th element of the l -th DPSS, \mathbf{s}_l , computed from equation (5). The first seven DPSS for $N_{\text{fft}} = 4096$ and $b_{\text{win}} = 4$ are shown in figure 1 and compared to the Hamming window, equation (9). The corresponding sample CSD matrix is

$$\mathbf{C}_j = \frac{1}{N_{\text{blk}}N_{\text{win}}} \sum_{k=1}^{N_{\text{blk}}} \sum_{l=1}^{N_{\text{win}}} \hat{\mathbf{q}}_j^{(k,l)} \left(\hat{\mathbf{q}}_j^{(k,l)} \right)^*. \quad (11)$$

Its matrix representation is also given by equation (3b), but for $M = N_{\text{blk}}N_{\text{win}}$ and $\hat{\mathbf{Q}}_j = [\hat{\mathbf{q}}_j^{(1,1)}, \hat{\mathbf{q}}_j^{(1,2)}, \dots, \hat{\mathbf{q}}_j^{(N_{\text{blk}}-1, N_{\text{win}})}, \hat{\mathbf{q}}_j^{(N_{\text{blk}}, N_{\text{win}})}]$.

The standard Welch SPOD algorithm is recovered if a single data window is
65 applied to multiple blocks, ($N_{\text{win}} = 1$, $N_{\text{blk}} > 1$); a multitaper SPOD algorithm if multiple orthogonal windows are applied to a single segment ($N_{\text{win}} > 1$, $N_{\text{blk}} = 1$); and a hybrid multitaper-Welch SPOD algorithm if blocking and orthogonal windowing are combined ($N_{\text{win}} > 1$, $N_{\text{blk}} > 1$).

2.4. Data

Example	Variables	N_x	$N_{y,r}$	N_t	Δt
Jet LES[35]	p	261	58	10000	0.2
Cavity PIV[36]	u, v	78	28	16000	$6.24 \cdot 10^{-5}$

Table 1: Database and spectral estimation parameters. All quantities are non-dimensionalized as described in the text.

70 Two sets of data are taken as examples. The first data set consists of 10000 snapshots of the pressure field from the wall-modeled large-eddy simulation (LES) of a turbulent $M = 0.9$ jet by Brès et al. [35]. This data has been studied

in detail using SPOD by Schmidt et al. [6]. The second data set consists of
 16000 snapshots of the streamwise and wall-normal velocity fields of the tur-
 75 bulent $M = 0.6$ flow over an open cavity acquired using time-resolved particle
 image velocimetry (TR-PIV) by Zhang et al. [36]. The open cavity has a length-
 to-depth ratio of $L/D = 6$ and a width-to-depth ratio of $W/D = 3.85$. More
 details on the experimental setup can be found in Zhang et al. [37], and an anal-
 ysis of the data, including using SPOD, in Zhang et al. [36]. All quantities are
 80 non-dimensionalized. For the jet LES data, the pressure is non-dimensionalized
 by twice the dynamic pressure, $\rho_j U_j^2$, where U_j and ρ_j are the jet exit velocity
 and density, respectively. Time is non-dimensionalized by D/c_∞ , where c_∞ is
 the far-field speed of sound and D the jet diameter. The dimensionless fre-
 quency can therefore be interpreted as the Strouhal number. The radial and
 85 streamwise coordinates are likewise given as multiples of D . The sampling time
 step Δt corresponds to 200 time steps of the original LES. The 16000 snapshots
 of the cavity PIV data were recorded over a period of $T = 1$ s. Using this period
 to non-dimensionalize time conveniently yields $\Delta f = 1$ and the non-dimensional
 frequency corresponds to the frequency in Hz. To facilitate the large parameter
 90 sweeps conducted for demonstration purposes in this study, the spatial resolu-
 tions of both databases were reduced by skipping every other grid point in the
 streamwise and normal directions. This reduction is accounted for in table 1
 and enables, in particular, the memory intensive computation of estimates that
 use a large number of tapers. We will later demonstrate how the combination of
 95 multitapering and segmenting can be used to reduce this computational burden
 for large data.

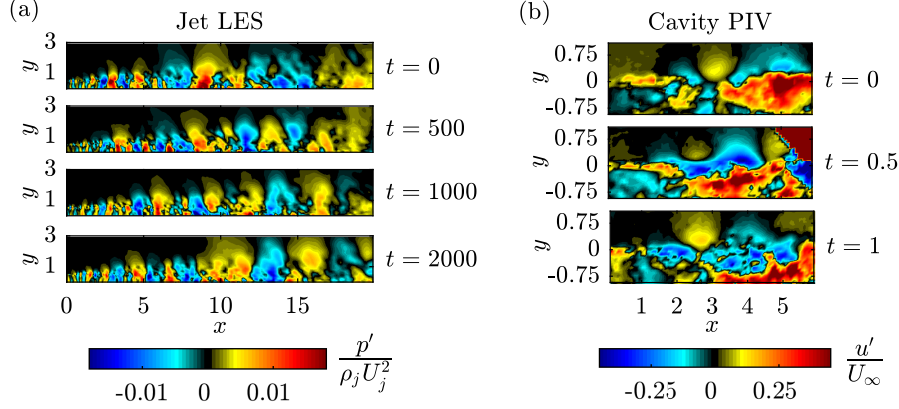


Figure 2: Instantaneous flow field visualizations: (a) fluctuating pressure of the turbulent jet LES; (b) streamwise fluctuating velocity of the open cavity PIV data.

Instantaneous flow field visualizations of both data sets are shown in figure 2, and relevant parameters are reported in table 1. The turbulent jet data are representative of turbulent flows with broadband turbulence spectra [35].
 100 The open cavity data are representative of turbulent flows with tonal peaks and an underlying broadband spectrum [36]. The two sets also differ in that measurement noise is present only in the experimental cavity flow data.

3. Results

For the remainder of this paper, we follow best spectral estimation practices
 105 [27], and set N_{ovlp} to 50% of N_{fft} for all Welch estimates. To compare multiple spectra for different parameters, we first focus on the leading eigenvalues, $\lambda_f^{(1)}$. The variance of the remaining eigenvalues is similar to that of the leading eigenvalue. We do, however, inspect the statistical convergence of the leading two SPOD modes, $\phi_j^{(1)}$ and $\phi_j^{(2)}$, and note that the convergence of the second mode
 110 is indicative of the convergence of the remaining modes. The leading SPOD modes play a special role in that they often capture physical mechanisms like resonances or hydrodynamic instabilities that dominate the dynamics at a given

frequency. The presence of such mechanisms leads to a rapid convergence of the leading mode and separation of the leading two eigenvalues that represent the mode energies. The latter phenomenon is often referred to as ‘low-rank behavior’, which should be understood in a physical rather than strictly mathematical sense ($\mathbf{C}_j \in \mathbb{R}^{N \times N}$ is positive semi-definite in general and has full rank only if constructed from $M = N$ independent realizations $\hat{\mathbf{q}}_j^{(k)}$, $k = 1, \dots, N$). We begin by comparing the basic Welch and multitaper estimates of the SPOD for the open cavity data in §3.1. In this context, we discuss the bias-variance tradeoff and some subtleties of the estimation process that are best explained by example. A comparative study of both test databases using multitaper-Welch estimates is presented in §3.2.

3.1. Standard Welch and single-block multitaper estimates of the SPOD

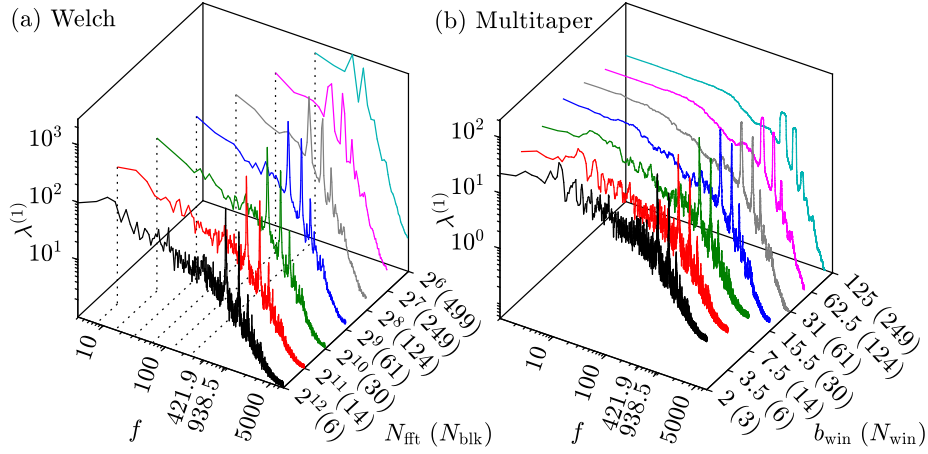


Figure 3: Leading SPOD mode spectra of the open cavity PIV data using different estimators: (a) Welch; (b) multitaper. The leading SPOD modes at $f = 421.9$ and $f = 938.5$ are shown in figure 4 below. Dashed lines in (a) indicate the lowest non-zero frequency for each N_{fft} .

Figure 3 compares SPOD eigenvalue spectra for the open cavity flow PIV data obtained using standard Welch and single-segment multitaper estimates. The free parameter that is varied for the Welch estimates in figure 3(a) is

number of samples per block, N_{fft} . For a fixed overlap of 50%, this also fixes the number of blocks, N_{blk} . For consistency, the Welch estimates are computed with the zeroth DPSS as the data window, and it was confirmed that the Welch estimates of the SPOD using \mathbf{s}_0 and a standard Hamming window, both shown in figure 1, are fundamentally similar. The multitaper estimates in figure 3(b) are obtained for the entire dataset of $N_{\text{fft}} = N_t = 16000$ snapshots, and the time-halfbandwidth product, b_{win} , was varied. The number of tapers used for each estimate, N_{win} , then directly follows from equation (7). The modes corresponding to two particular frequencies of $f = 938.5$ and $f = 421.9$ that are labeled on the frequency axis will be discussed later. The higher frequency corresponds to the highest peak in the spectrum and is associated with the first Rossiter mode [38], a resonant hydrodynamic instability. The lower frequency corresponds to the much broader peak in the region of high variance that appears close to the Rossiter frequency on the logarithmic scale.

Welch’s method is asymptotically unbiased for stationary data and converges to the true spectral density as $N_{\text{blk}}, N_{\text{fft}} \rightarrow \infty$. For limited data, however, care has to be taken regarding the balance of bias and variance. This becomes clear in figure 3(a). For $N_{\text{fft}} = 2^{12} = 4096$ (black line), the data are segmented into only six blocks, and the noisy appearance of the SPOD spectrum indicates the high variance of the estimate. As N_{fft} is decreased, the number of blocks increases, and the variance is significantly reduced. Observe, however, that the distinct spectral peaks at the Rossiter and some higher frequencies vanish simultaneously. This is evidence of bias due to the decreased frequency resolution, $\Delta f_{\text{Welch}} = \frac{1}{N_{\text{fft}} \Delta t} = \frac{f_s}{N_{\text{fft}}}$, with decreasing N_{fft} and the associated spectral leakage of the under-resolved features. Similar to Welch’s method, multitaper estimates of the spectral density based on DPSS are asymptotically unbiased and their mean square error converges to zero as $N \rightarrow \infty$ (but at the same time $\delta f \rightarrow 0$ with $b_{\text{win}} \rightarrow \infty$, and $N_{\text{win}} \leq 2b_{\text{win}}$ but $N_{\text{win}} \rightarrow \infty$; see [39] for details). As b_{win} is increased in figure 3(b), the same qualitative trends are observed as for the Welch estimator, i.e., the variance decreases as the bias increases. The multitaper estimate retains the same frequency resolution, or

bin size, of $\Delta f = \frac{1}{\Delta t N_t} = 1$, whereas for the Welch estimate, it increases from
160 $\Delta f_{\text{Welch}} = \frac{1}{\Delta t N_{\text{fft}}} = 3.9$ for $N_{\text{fft}} = 4096$ to $\Delta f_{\text{Welch}} = 250$ for $N_{\text{fft}} = 64$. For the
multitaper estimator at a fixed N , the time-halfbandwidth product, b_{win} , deter-
mines the frequency bandwidth, $\Delta f_{\text{DPSS}} = f_s \delta f$, over which the DPSS tapers
yield an integrated average. Welch and multitaper estimates may be compared
if they average over a similar frequency band, this is if $\Delta f_{\text{Welch}} \approx \Delta f_{\text{DPSS}}$.

165 For the Welch estimate with $N_{\text{fft}} = 512$, for example, the frequency res-
olution is reduced by a factor of $\frac{N_t}{N_{\text{fft}}} = 31.25$. Because $\Delta f = 1$ (16000
samples recorded at 16 kHz over 1 sec) in the experiment, the same value,
 $\Delta f_{\text{Welch}} = 31.25$, is obtained for the frequency bin size. Since b_{win} determines
the resolution bandwidth in multiples of Δf , we consider $b_{\text{win}} = 31$, and may
170 confirm that $\Delta f_{\text{DPSS}} = 31$ (again, because $\Delta f = 1$). We observe that the Welch
estimate for $N_{\text{fft}} = 512$ (blue line in figure 3(a)) and the multitaper estimate for
 $b_{\text{win}} = 31$ (gray line in figure 3(b)), are indeed qualitatively similar in terms of
their variance and resolution of peaks. If the variance is further reduced (cyan
lines in figure 3(a,b)), then the bias is further increased, and spectral peaks are
175 increasingly under-resolved for the Welch estimator, and flattened and broad-
ened for the multitaper estimator. An advantage of multitaper SPOD for com-
parable bias and variance is that it resolves lower frequencies down to Δf . A
disadvantage is that it becomes computationally expensive for large N_{win} , as
it requires parallel processing of N_{win} realizations of the Fourier transform of
180 the full data. The computational cost associated with multitaper estimates is
discussed in more detail in §4.

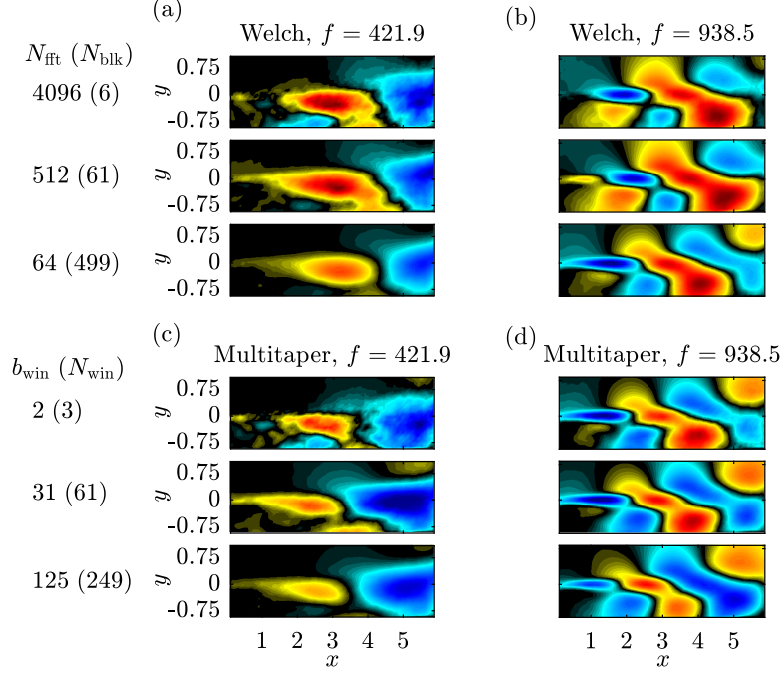


Figure 4: Real part of the leading SPOD modes of the open cavity PIV data using different estimators: (a,b) Welch for varying $N_{\text{fft}}, N_{\text{blk}}$; (c,d) multitaper for varying $b_{\text{win}}, N_{\text{win}}$. The leading SPOD modes at $f = 421.9$ and $f = 938.5$ are shown in the left and right columns, respectively (labelled in figure 3 above).

Figure 4 shows SPOD modes obtained using the Welch (top) and multi-
 taper (bottom) estimators at the two different frequencies, respectively repre-
 senting the low-frequency region with high variance (left) and the tonal peak
 frequency associated with the first Rossiter mode (right). The phases of the com-
 185 plex modes have been approximately aligned for comparability. At $f = 938.5$,
 the Rossiter mode in figure 4(b,d) is extracted by both approaches, and well-
 converged for all parameters. This is typical for hydrodynamic instabilities and
 resonant mechanisms that are both energetic and evolve coherently over large
 190 regions in space and time. At the lower frequency of $f = 421.9$, no such mech-
 anism is present, and both the Welch estimate for $N_{\text{blk}} = 6$ and the multitaper

estimate for $N_{\text{win}} = 3$ appear noisy. Well-converged modes are obtained when N_{fft} is decreased (and therefore N_{blk} increased), or b_{win} (and therefore N_{win}) are increased for the multitaper estimate.

195 3.2. Multitaper-Welch estimates of the SPOD

As demonstrated in §3.1 above, the multitaper method can, in principle, obtain low-variance estimates without segmentation if a large number of data windows is used. Doing so, however, quickly becomes computationally prohibitive. Furthermore, the estimate at the lowest resolvable frequency relies on
200 a single realization of the flow process if $N = N_t$, and is therefore not statistically representative. These issues are addressed by applying multiple orthogonal windows to overlapping segments, that is, the multitaper-Welch estimate given by equation (11) with $N_{\text{blk}}, N_{\text{win}} > 1$. Our results from §3.1 suggest that Welch and multitaper SPOD yield comparably well-converged modes and spectra for
205 similar bandwidths, $\Delta f_{\text{Welch}} \approx \Delta f_{\text{DPSS}}$, leaving computational cost as the only advantage of the Welch estimator.

For the following study of the performance of the multitaper-Welch estimator, we chose the highest N_{fft} possible that results in the lowest number of blocks that still gives us confidence in the statistical convergence at low frequencies.
210 That segment length is $N_{\text{fft}} = 4096$, resulting in $N_{\text{blk}} = 3$ and $N_{\text{win}} = 6$ for the jet and cavity data, respectively.




Figure 5: SPOD mode spectra of the cavity PIV data using different estimators with comparable bandwidth, Δf : (a) Welch with $N_{\text{fft}} = 512$; (b) multitaper with $b_{\text{win}} = 31$; (c) multitaper-Welch with $N_{\text{fft}} = 4096$, $b_{\text{win}} = 8$. Colors correspond to those in figures 3(a), 3(b), and 6(b) for $b_{\text{win}} = 9$, respectively. Red lines show the integral power spectral density (sum of eigenvalues at each frequency) and dotted lines the lowest non-zero frequencies of all three estimates.

Before systematically assessing the effect of the bandwidth parameter, we start with a three-way comparison, including the hybrid multitaper-Welch estimator, at comparable variance. Following the discussion in §3.1, a comparable

215 variance is achieved by adjusting the spectral estimation parameters such that $\Delta f_{\text{Welch}} \approx \Delta f_{\text{DPSS}} \approx \Delta f_{\text{M-W}}$, where $\Delta f_{\text{M-W}}$ is the bandwidth of the hybrid estimator. Figure 5(a) and 5(b) show the full SPOD spectra for the Welch and multitaper estimators with $\Delta f \approx 31$, previously shown and discussed in the

context of figure 3(a) and 3(b). The same colors as in figure 3 are used. Noting
220 that the sum of all eigenvalues corresponds to the energy density, the spectra
are multiplied by the number of frequency bins to collapse them on the same
scale. The bandwidth of the multitaper-Welch method, $\Delta f_{\text{M-W}} = \frac{1}{N_{\text{fft}} \Delta t} \frac{\delta f}{\Delta t}$,
results from the combined effects of windowing and tapering. For a fixed seg-
ment length of $N_{\text{fft}} = 4096$, we set $\Delta f_{\text{M-W}} = 31.25$ by choosing $b_{\text{win}} = 8$.
225 The corresponding SPOD spectrum is shown in green in 5(c) (same color as
for the closest bandwidth parameter of $b_{\text{win}} = 9$ in figure 6 below). We may
convince ourselves that the variances of all three estimates in figure 5(a,b,c) are
indeed visually very similar. The computational cost of the different estimators
is discussed in detail in §4 below. Of the three estimates shown here, the Welch
230 estimate computed that fastest, the Multitaper-Welch estimate took ~ 15 times,
and the Multitaper-only estimate ~ 45 times as long. The main advantage of the
multitaper and hybrid estimators is their higher resolution, in particular at the
lower frequencies that contain the largest fraction of the integral energy in most
flows. The lowest resolvable frequency for the Welch estimator is $f = 31.25$,
235 whereas it is one for the Multitaper method.

Next, we vary the bandwidth parameter, b_{win} , and therefore the number of
Slepian tapers, N_{win} , to systematically study the bandwidth effect.

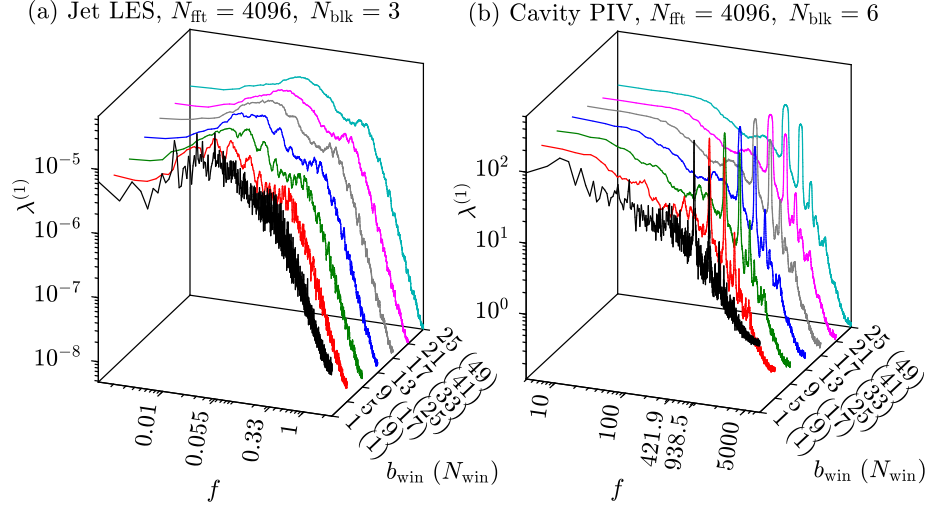


Figure 6: Leading SPOD mode spectra using the multitaper-Welch estimator: (a) turbulent jet LES; (b) open cavity PIV data. The leading SPOD modes at $f = 0.055$ and $f = 0.33$. The segment length is $N_{\text{fft}} = 4096$, resulting in $N_{\text{blk}} = 3$ and $N_{\text{blk}} = 6$ for the jet and cavity data, respectively.

Figure 6 shows the SPOD spectra for the leading modes and both cases. The bandwidth parameter is varied from $b_{\text{win}} = 1$ to 25, resulting in a number of windows between $N_{\text{win}} = 1$ (Welch estimator) and 49. The total number of
240 realizations, $N_{\text{blk}}N_{\text{fft}}$, is between 3 and 147 for the jet data and between 6 and 294 for the cavity data. A similar decrease in variance is observed for both data as the bandwidth is increased. This decrease is linked to an increase in bias that leads to the loss of certain features, such as several broader peaks at low
245 frequencies in the jet data (still resolved for $b_{\text{win}} = 9$, green line) and several double peaks at high frequencies in the cavity data (still clearly visible for $b_{\text{win}} = 17$, gray line). Under the caveat that the bias cannot be quantified because the true spectrum is unknown, these observations reflect the usual variance-bias tradeoff.

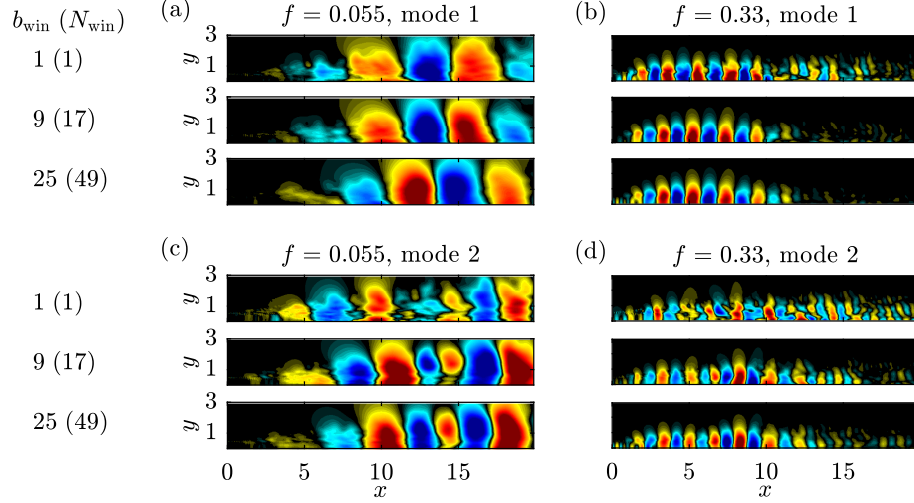


Figure 7: Two leading SPOD modes for the jet LES data using multitaper-Welch estimates of different bandwidths: (a) first mode for $f = 0.055$; (b) first mode for $f = 0.33$; (c) second mode for $f = 0.055$; (d) second mode for $f = 0.33$. The corresponding leading SPOD energy spectra are reported in figure 6.

250 The first and second SPOD modes of the jet data at the two representative frequencies are shown in figure 7. Modes obtained for a single, an intermediate, and the highest number of tapers are compared. As expected, better-converged modes are obtained for higher bandwidths. Take as an example the leading SPOD mode for $f = 0.33$ shown in 7(b). For $b_{\text{win}} = 1$ and a single taper, the
255 mode appears noisy, particularly in the downstream region for $x \gtrsim 10$. For the two higher bandwidths, the waveform is much more compact and smooth, and very similar for $b_{\text{win}} = 9$ and 25. This is indirect evidence of the statistical convergence of the mode. For the second (first subdominant) SPOD mode shown in figure 7(d), continual improvement is observed as the bandwidth is increased
260 and a double-lobed structure emerges for the highest value of $b_{\text{win}} = 25$.

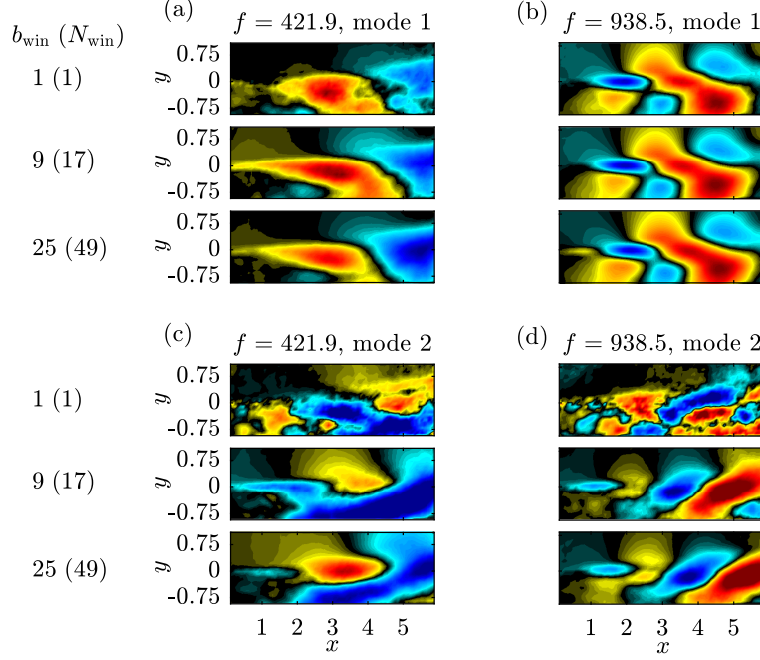


Figure 8: Same as figure 7 but for the cavity PIV data: (a) first mode for $f = 421.9$; (b) first mode for $f = 938.5$; (c) second mode for $f = 421.9$; (d) second mode for $f = 938.5$.

Similar observations are made for the SPOD modes of the cavity PIV data shown in figure 8. The improvement of mode convergence with increasing bandwidth is apparent, in particular for the second SPOD modes in figure 8(c,d). A clear spatial pattern that is physically interpretable and quantifiable in terms of, for example, streamwise and wall-normal wavenumbers, is only revealed for the two higher bandwidths. A notable exception, as discussed earlier in the context of figure 4, is the Rossiter mode shown in figure 8(b) at a frequency of $f = 938.5$. It is associated with the overall largest SPOD eigenvalue and is so prevalent that it can be spotted in the instantaneous velocity fields in figure 2(b).




Figure 9: Multitaper-Welch SPOD spectra for the cavity PIV data. The data are segmented into $N_{\text{blk}} = 6$ blocks of length $N_{\text{fft}} = 4096$ and $b_{\text{win}}, N_{\text{win}}$ are varied: (a) $b_{\text{win}} = 1, N_{\text{win}} = 1$; (b) $b_{\text{win}} = 9, N_{\text{win}} = 17$; (c) $b_{\text{win}} = 17, N_{\text{win}} = 33$; (d) $b_{\text{win}} = 25, N_{\text{win}} = 49$.

Finally, we investigate in figure 9 the full SPOD spectrum for the cavity PIV data and the relative energy content of the leading and higher SPOD modes. Four values of b_{win} between 1, in figure 9(a), and the maximum of 25, in 9(d), are considered. It is, as expected, observed that the variance of the higher mode spectra decreases in a very similar manner to that of the leading mode. The normalized cumulative energy, that is, the cumulative sum of the eigenvalues from the highest to the lowest at each frequency, is indicated by the color. The number of modes required to capture 90% of the energy at any given frequency corresponds to the number of all colored (non-gray) lines. For the two highest peaks and the lowest bandwidth of $b_{\text{win}} = 1$ shown in 9(a), the first mode alone account for more than 90% of the energy. As the bandwidth is increased,

the peaks flatten, and this number reduces to 60%-80%. At the same time, the number of modes required to reach 90% stabilizes, and the main effect of increasing b_{win} is, apart from the discussed effects on bias and variance, that more eigenvalues of lower energy are added to the bottom of the spectrum. This, again, is an indicator of statistical convergence.

4. Computational performance

The computational cost of the SPOD algorithm [5, 28] implemented in the open-source Matlab code **SPOD**, and extended to multitaper estimators in the present work, is dominated by two main loops (see Schmidt and Towne [40] for a detailed description and illustration of the algorithm). The first loop goes over all N_{blk} blocks and computes, according to equation (10), the temporal DFT of each block for each taper. The second loop goes over all N_{fft} frequencies. For each frequency, the CSD matrix, as defined in equation (4b), is assembled and its eigenvalue decomposition (EVD), equation (4a), is computed. Most of the compute time is, for most cases, spent in this second loop. The only exceptions are the multitaper estimates with the two smallest numbers of tapers ($N_{\text{win}} = 3$ and 6), for which the DFTs take most of the time. Clearly, any performance advantage of the multitaper-Welch approach comes at the additional computational cost associated with processing a factor equal to the number of tapers of additional Fourier transforms.

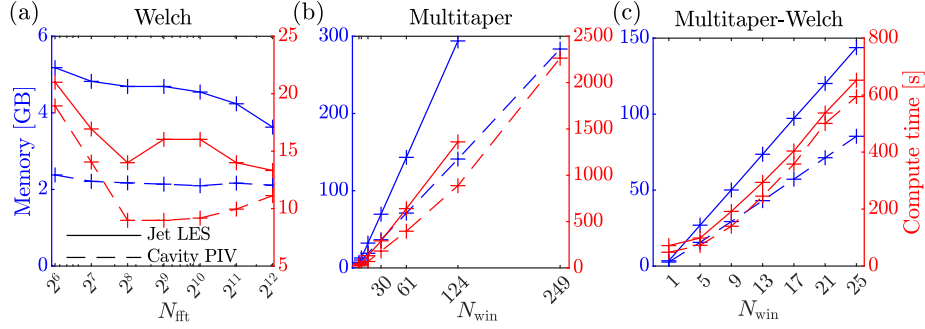


Figure 10: Computational performance in terms of memory usage (blue) and compute time (red): (a) Welch; (b) multitaper; (c) multitaper-Welch. Both the jet LES (solid lines) and cavity PIV (dashed lines) data are considered. Refer to figures 3 and 6 for the corresponding SPOD spectra. All axes but the abscissa in (a) are on a linear scale.

Figure 10 reports the memory usage and compute times for the Welch and multitaper estimates from §3.1, and the multitaper-Welch estimates from §3.2. The computations were conducted on a workstation with 20 Intel Xeon cores
305 and 512 GB of memory. The memory consumption and compute times of the Welch estimates shown in figure 10(a) vary non-monotonically with the segment length, and do not significantly change as N_{fft} is varied from 64 to 4096. This can be explained by the balance between the number of segments and their size that is inherent to the algorithm. If N_{fft} is large, then a small number
310 of $N_{\text{blk}}N_{\text{dof}}$ DFTs of length N_{fft} have to be computed in the first loop, and a large number of N_{fft} small EVDs of size $N_{\text{blk}} \times N_{\text{blk}}$ have to be solved in the second. These trends are reversed for small N_{fft} , which results in large N_{blk} , and the overall computational cost is comparable. A very different behavior is observed for the multitaper and multitaper-Welch estimates in figure 10(b)
315 and 10(c), respectively. The use of multiple tapers proportionally inflates the number of DFTs that have to be stored in memory and processed from N_{blk} to $N_{\text{blk}}N_{\text{win}}$ ($N_{\text{blk}} = 1$ for multitaper-only), independent of their length. This proportional increase of the the problem size with increasing N_{win} is directly reflected in the near-linear scaling of memory and compute time for estimates

320 that use multiple tapers. Note, in particular, that the multitaper-only estimates
 with the largest values of N_{win} require almost 300 GB of memory for both data
 sets. The data reduction described in §2.4 was largely motivated by the desire
 to include these extreme examples in this study for demonstration purposes.
 As demonstrated in §3.1, very large values of N_{win} lead to a large bias, and
 325 are therefore not recommended in any case. The multitaper-Welch estimates
 shown and discussed in the context of figure 6 provide more balanced results
 at significantly lower computational cost. The lower computational cost of the
 hybrid approach becomes apparent from comparing the memory and compute
 times of the multitaper estimates in panel 10(b) to those of the multitaper-
 330 Welch algorithm in panel 10(c). This assessment of computational performance
 is taken into account for the best practice recommendations presented next as
 part of the discussion in §5.

5. Discussion

In the light of the above comparative study, we start by assessing the per-
 335 formance of multitaper and multitaper-Welch estimates of the SPOD compared
 to the standard Welch approach. By comparing SPOD spectra obtained us-
 ing Welch-only and multitaper-only estimators in figure 3, we have established
 that both algorithms are comparable in terms of variance, bias, and mode con-
 vergence if the frequency resolution of the Welch estimate is similar to the
 340 bandwidth of the multitaper estimate. This leaves the multitaper estimate with
 the advantage of a higher resolution and the associated ability to resolve lower
 frequencies. Hybrid multitaper-Welch estimates offer the flexibility to control to
 a large degree the resolution, bias, and variance. In the limit of unit bandwidth
 and a single taper, the classical Welch estimator is recovered. This allows us
 345 to assess the performance advantage of multitaper-Welch estimates over Welch
 estimates directly from figures 6-9, that all contain this limiting case. All com-
 parisons show that additional tapers reduce the variance of the basic Welch
 estimate and lead to better-converged modes. Similar to the findings of Bronez

[30] on the estimation of power spectra from time signals, our results show that
 350 the multitaper-Welch approach generally outperforms the Welch-only approach
 in terms of variance at fixed resolution, and in terms of resolution if the variance
 is comparable. This performance advantage comes at the additional computa-
 tional cost associated with processing a factor equal to the number of tapers of
 additional Fourier transforms.

355 We hence recommend using the multitaper-Welch estimator to compute the
 SPOD and the following best practice: for given data consisting of N_t snap-
 shots, first select N_{fft} to obtain the smallest small number of blocks that still
 ensures that the dynamics at the lowest frequencies are statistically adequately
 represented. Then choose either the minimum bandwidth necessary to reduce
 360 the variance (or increase the mode convergence) to the desired level, *or* select
 the maximum bandwidth acceptable in terms of bias (or any value in between
 the two). In settings where well-separated blocks of data are obtained individ-
 ually, for example, in carefully designed experiments like those by Citriniti and
 George [41], multitaper estimates from each block can be leveraged to reduce
 365 the number of required blocks, or reduce the variance and improve the mode
 convergence for an existing set of data.

While SPOD computes modes that optimally represent the data in terms
 of its second-order statistics, i.e., variance or energy, multitaper estimates are
 equally applicable to estimate higher-order statistics. They can hence be lever-
 370 aged to improve estimates of the bispectral mode decomposition [42]. SPOD can
 quickly become computationally intractable for very large data. This problem
 can be addressed by updating algorithms that converge the SPOD from stream-
 ing data that become available one snapshot at a time. While impossible for the
 multitaper-only estimator, the algorithm proposed by Schmidt and Towne [40]
 375 can readily be extended to accommodate multiple data windows and compute
 multitaper-Welch estimates on the fly.

The multitaper-Welch estimator using DPSS was implemented in the exist-
 ing open-source Matlab code SPOD [5]. The second argument of the function
 can be used to specify N_{fft} (NFFT) and b_{win} (BW) using the syntax `[L,P,F] =`

380 $\text{SPOD}(\mathbf{X}, [\text{NFFT BW}], \dots)$. The number of DPSS tapers is then determined by equation (7).

Acknowledgments. OTS gratefully acknowledges support from Office of Naval Research grant N00014-20-1-2311 and NSF grant CBET 2046311. I would like to thank Lou Cattafesta and Yang Zhang for providing the TR-PIV data, Gregg
385 Abate for connecting me to Lou, Guillaume Brès for generating the turbulent jet database, and Eduardo Martini and Akhil Nekkanti for their constructive feedback. The TR-PIV data was created with support from AFOSR Award Number FA9550-17-1-0380. Creation of the LES data was supported by NAVAIR SBIR under the supervision of J. T. Spyropoulos, with computational resources provided by DoD HPCMP at the ERDC DSRC supercomputer facility.
390

Data and code availability. The datasets analyzed in this study are currently not publicly available as they were generated by other researchers (see §2.4, table 1) as part of DoD funded independent studies (see *Acknowledgments*), but are available from the corresponding author on reasonable request. A reduced
395 version of the turbulent jet data set is publicly available as part of the open-source Matlab code `SPOD`, that now supports multitaper-Welch estimates for `SPOD`.

References

1. Lumley JL. Stochastic tools in turbulence. New York: Academic Press;
400 1970.
2. Glauser MN, Leib SJ, George WK. Coherent structures in the axisymmetric turbulent jet mixing layer. *Turbulent shear flows* 1987;5:134–45.
3. Schmid PJ. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics* 2010;656:5–28. doi:10.1017/S0022112010001217.
405

4. McKeon BJ, Sharma AS. A critical-layer framework for turbulent pipe flow. *Journal of Fluid Mechanics* 2010;658:336–382. doi:10.1017/S002211201000176X.
5. Towne A, Schmidt OT, Colonius T. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics* 2018;847:821–867. doi:10.1017/jfm.2018.283.
6. Schmidt OT, Towne A, Rigas G, Colonius T, Brès GA. Spectral analysis of jet turbulence. *Journal of Fluid Mechanics* 2018;855:953–982. doi:10.1017/jfm.2018.675.
7. Arndt REA, Long DF, Glauser MN. The proper orthogonal decomposition of pressure fluctuations surrounding a turbulent jet. *Journal of Fluid Mechanics* 1997;340:1–33.
8. Citriniti JH, George WK. Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition. *Journal of Fluid Mechanics* 2000;418:137–66.
9. Nidhan S, Chongsiripinyo K, Schmidt OT, Sarkar S. Spectral proper orthogonal decomposition analysis of the turbulent wake of a disk at $Re=50000$. *Physical Review Fluids* 2020;5(12):124606. doi:10.1103/PhysRevFluids.5.124606.
10. Gordeyev SV, Thomas FO. Coherent structure in the turbulent planar jet. part 1. extraction of proper orthogonal decomposition eigenmodes and their self-similarity. *Journal of Fluid Mechanics* 2000;414:145–94.
11. Gudmundsson K, Colonius T. Instability wave models for the near-field fluctuations of turbulent jets. *Journal of Fluid Mechanics* 2011;689:97–128.
12. Hellström LHO, Smits AJ. The energetic motions in turbulent pipe flow. *Physics of Fluids* 2014;26(12):125102.

13. Tutkun M, George WK. Lumley decomposition of turbulent boundary layer
435 at high reynolds numbers. *Physics of Fluids* 2017;29(2):020707.
14. Araya DB, Colonius T, Dabiri JO. Transition to bluff-body dynamics
in the wake of vertical-axis wind turbines. *Journal of Fluid Mechanics*
2017;813:346–81.
15. Abreu LI, Cavalieri AVG, Wolf W. Coherent hydrodynamic waves and
trailing-edge noise. In: *23rd AIAA/CEAS aeroacoustics conference*; vol.
440 AIAA 2017-3173. 2017;doi:10.2514/6.2017-3173.
16. He X, Fang Z, Rigas G, Vahdati M. Spectral proper orthogonal decomposi-
tion of compressor tip leakage flow. *Physics of Fluids* 2021;33(10):105105.
17. Li XB, Chen G, Liang XF, Liu DR, Xiong XH. Research on spectral estima-
445 tion parameters for application of spectral proper orthogonal decomposition
in train wake flows. *Physics of Fluids* 2021;33(12):125103.
18. Schmidt OT, Mengaldo G, Balsamo G, Wedi NP. Spectral em-
pirical orthogonal function analysis of weather and climate data.
Monthly Weather Review 2019;147(8):2979–95. URL: <https://doi.org/10.1175/MWR-D-18-0337.1>.
450 [doi:10.1175/MWR-D-18-0337.1](https://doi.org/10.1175/MWR-D-18-0337.1).
[arXiv:https://doi.org/10.1175/MWR-D-18-0337.1](https://doi.org/10.1175/MWR-D-18-0337.1).
19. Sanjose M, Towne A, Jaiswal P, Moreau S, Lele S, Mann A. Modal anal-
ysis of the laminar boundary layer instability and tonal noise of an air-
foil at reynolds number 150,000. *International Journal of Aeroacoustics*
455 2019;18(2-3):317–50.
20. Nekkanti A, Schmidt OT. Modal analysis of acoustic direc-
tivity in turbulent jets. *AIAA Journal* 2020;0(0):1–12. URL:
<https://doi.org/10.2514/1.J059425>. doi:10.2514/1.J059425.
[arXiv:https://doi.org/10.2514/1.J059425](https://doi.org/10.2514/1.J059425).
- 460 21. Baars WJ, Tinney CE. Proper orthogonal decomposition-based spectral
higher-order stochastic estimation. *Physics of Fluids* 2014;26(5):055112.

22. Nekkanti A, Schmidt OT. Frequency–time analysis, low-rank reconstruction and denoising of turbulent flows using spod. *Journal of Fluid Mechanics* 2021;926:A26. doi:10.1017/jfm.2021.681.
- 465 23. Ghate AS, Towne A, Lele SK. Broadband reconstruction of inhomogeneous turbulence using spectral proper orthogonal decomposition and gabor modes. *Journal of Fluid Mechanics* 2020;888:R1. doi:10.1017/jfm.2020.78.
- 470 24. Chu T, Schmidt OT. A stochastic spod-galerkin model for broadband turbulent flows. *Theoretical and Computational Fluid Dynamics* 2021;URL: <https://doi.org/10.1007/s00162-021-00588-6>. doi:10.1007/s00162-021-00588-6.
- 475 25. Towne A. Space-time galerkin projection via spectral proper orthogonal decomposition and resolvent modes. In: *AIAA Scitech 2021*; vol. AIAA 2021-1676. 2021:1676.
26. Pain R, Weiss PE, Deck S, Robinet JC. Large scale dynamics of a high reynolds number axisymmetric separating/reattaching flow. *Physics of Fluids* 2019;31(12):125119.
- 480 27. Welch P. The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics* 1967;15(2):70–3.
28. Schmidt OT, Colonius T. Guide to spectral proper orthogonal decomposition. *AIAA Journal* 2020;58(3):1023–33. doi:10.2514/1.J058809.
- 485 29. Thomson DJ. Spectrum estimation and harmonic analysis. *Proceedings of the IEEE* 1982;70(9):1055–96.
30. Bronez TP. On the performance advantage of multitaper spectral analysis. *IEEE Transactions on Signal Processing* 1992;40(12):2941–6.

31. Geoga CJ, Haley CL, Siegel AR, Anitescu M. Frequency–wavenumber spectral analysis of spatio-temporal flows. *Journal of Fluid Mechanics* 2018;848:545–559. doi:10.1017/jfm.2018.366.
32. Slepian D, Pollak HO. Prolate spheroidal wave functions, fourier analysis and uncertainty—i. *Bell System Technical Journal* 1961;40(1):43–63.
33. Slepian D. Prolate spheroidal wave functions, fourier analysis, and uncertainty—v: The discrete case. *Bell System Technical Journal* 1978;57(5):1371–430.
34. Heinzel G, Rüdiger A, Schilling R. Spectrum and spectral density estimation by the discrete fourier transform (dft), including a comprehensive list of window functions and some new at-top windows 2002;.
35. Brès G, Jordan P, Le Rallic M, Jaunet V, Cavalieri AVG, Towne A, Lele S, Colonius T, Schmidt OT. Importance of the nozzle-exit boundary-layer state in subsonic turbulent jets. *Journal of Fluid Mechanics* 2018;851:83–124. doi:10.1017/jfm.2018.476.
36. Zhang Y, Cattafesta L, Ukeiley L. A spectral analysis modal method applied to cavity flow oscillations. In: *TSFP11, Southampton, UK*. 2019;.
37. Zhang Y, Cattafesta L, Ukeiley L. Identification of coherent structures in cavity flows using stochastic estimation and dynamic mode decomposition. In: *TSFP10, Chicago, USA*. 2017: 3.
38. Rossiter JE. Wind tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. *RAE Technical Report No 64037* 1964;.
39. Lii KS, Rosenblatt M. Prolate spheroidal spectral estimates. *Statistics & Probability Letters* 2008;78(11):1339–48.
40. Schmidt OT, Towne A. An efficient streaming algorithm for spectral proper orthogonal decomposition. *Computer Physics Communications* 2018;URL: <http://www.sciencedirect.com/science/article/pii/>

- 515 S0010465518304016. doi:<https://doi.org/10.1016/j.cpc.2018.11.009>.
41. Citriniti JH, George WK. Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition. *Journal of Fluid Mechanics* 2000;418:137–66.
- 520 42. Schmidt OT. Bispectral mode decomposition of nonlinear flows. *Nonlinear Dynamics* 2020;(102(4)):2479–501. doi:10.1007/s11071-020-06037-z.