

Band-ensemble spectral proper orthogonal decomposition with frequency attribution

Jakob G. R. von Saldern ^{*}

Department of Civil and Mechanical Engineering,

California Institute of Technology, Pasadena, California 91125, USA

and Laboratory for Flow Instabilities and Dynamics, Technische Universität Berlin, 10623 Berlin, Germany

Oliver T. Schmidt 

Department of Mechanical and Aerospace Engineering,

University of California, San Diego, La Jolla, California 92093, USA

Philipp Godbersen 

Institute of Aerodynamics and Flow Technology,

German Aerospace Center (DLR), Göttingen 37073, Germany

J. Moritz Reumschüssel 

Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom

and Chair of Fluid Dynamics, Technische Universität Berlin, 10623 Berlin, Germany

Tim Colonius 

Department of Civil and Mechanical Engineering,

California Institute of Technology, Pasadena, California 91125, USA



(Received 6 February 2026; accepted 15 September 2026; published 23 October 2026)

This study presents band-ensemble spectral proper orthogonal decomposition (bSPOD). The approach is inspired by frequency smoothing, a method used to reduce estimator variance in power spectral density estimates, and is here extended to SPOD. The algorithm estimates SPOD modes from consecutive Fourier coefficients obtained from a single Fourier transform of the full-time record and thus avoids time segmentation. In this study, bSPOD is applied to artificial test data and to a particle image velocimetry data set of a broadband-tonal cavity flow. Compared to the more commonly used Welch-based SPOD formulation, bSPOD reduces spectral leakage, permits increased frequency resolution, and retains frequency information of tonal components at comparable computational cost. These features enable reduced estimator variance while maintaining low bias for tonal components, making bSPOD particularly effective for broadband-tonal flows.

DOI: [10.1103/b9hq-xvpx](https://doi.org/10.1103/b9hq-xvpx)

^{*}Contact author: j.vonsaldern@tu-berlin.de

I. INTRODUCTION

Turbulent flow fields are inherently high-dimensional and chaotic. Identifying organized spatiotemporal patterns within these chaotic dynamics, has long been central in fluid mechanics [1–6]. These so-called coherent structures are typically associated with high energy fluctuations and often underlie undesirable phenomena such as sound generation [7,8], stall cell formation on airfoils [9], combustion unsteadiness [10,11], and structural vibration [12].

Techniques for processing such chaotic data have evolved over the last century. A central tool in this context is the power spectral density (PSD) which describes how a signal’s power is distributed over temporal frequency. In fluid mechanics, PSDs are widely used for processing velocity-probe and pressure-tap signals. The PSD is the Fourier transform of the autocorrelation function but is typically estimated from periodograms [13,14]. Because turbulence is chaotic, averaging is used to reduce estimator variance. For an estimate of the PSD, $\hat{S}(\omega)$, this variance can be written as

$$\text{Var}[\hat{S}(\omega)] = \mathbb{E}[(\hat{S}(\omega) - \mathbb{E}[\hat{S}(\omega)])^2], \quad (1)$$

and quantifies the deviation of a finite-sample spectral estimate from its theoretical value. Since the latter is generally unknown *a priori*, the estimation variance cannot usually be evaluated directly. However, the estimation variance decreases with increasing number of statistically independent realizations employed in the averaging procedure. In practice, a high estimation variance manifests itself as noisy, jittering, or irregular spectral estimates. Consequently, the term estimation variance is often used more loosely to describe these statistical fluctuations, and is used in this sense throughout the present work. Among averaging methods, Bartlett’s [15] approach is the most commonly used: assuming ergodicity, the method divides the record into segments, computes a PSD for each, and averages the results. This reduces estimation variance at the cost of decreased frequency resolution, which is set by the segment length. Welch [16] extended this approach by introducing overlapping segments and applying a taper to each segment to reduce spectral leakage.

Thomson [17] introduced the multitaper estimator, which multiplies the full-time record by a set of orthogonal tapers to obtain approximately independent realizations. These tapered signals are Fourier transformed and averaged across tapers. This yields a PSD on the same frequency grid as the full-record DFT, but with a lower effective resolution because the estimate equals the true spectrum convolved with a spectral window formed from the tapers’ Fourier transforms [18]. At matched effective resolution, multitaper generally reduces spectral leakage and provides comparable or lower variance than segment-based estimates under standard Gaussian assumptions [19]. The tradeoff is computational cost, since a full-record transform is computed for each taper.

A related approach to variance reduction is frequency smoothing [20–22]. Daniell [20] first proposed averaging periodogram values over adjacent frequencies, and Bartlett [21] noted the connection to truncating the autocorrelation. Blackman and Tukey [22] then formalized the lag-window autocorrelation estimator by tapering the autocorrelation at large lags. The frequency-smoothing and lag-window views are equivalent when the smoothing filter and the lag window are a Fourier pair. In simple terms, this approach reduces estimator variance by averaging a full-record periodogram over neighboring frequencies, at the cost of lower effective frequency resolution.

Time segmentation, multitapering and frequency smoothing all reduce estimator variance but increase bias, i.e., the deviation between the expected spectral estimate and the true spectrum, because finite-window operations either in time or frequency domain result in a spread of power across neighboring frequency bins. Where the PSD varies rapidly, for example at tonal peaks in the spectrum, stronger averaging broadens the narrow spectral features. A key distinction between the three approaches is that time-segmentation first splits the record into shorter segments, computes periodograms for each, and averages them; the effective frequency resolution is set by the segment length and is essentially uniform across the spectrum. In contrast, multitapering allows the number of tapers and their weights, which together determine the effective frequency resolution, to vary with frequency. Similarly, in frequency smoothing a single full-record periodogram is computed and then

averaged across frequency, so the smoothing filter and its bandwidth can be chosen frequency by frequency, enabling a local bias-variance tradeoff.

In fluid mechanics, not only temporal but also spatially resolved data are often available, for example, from particle image velocimetry (PIV) or from arrays of pressure or velocity probes sampled simultaneously. Consequently, in addition to temporal statistics one can evaluate spatial statistics to extract coherent structures. Spatial structures can be identified with proper orthogonal decomposition (POD), which, via an eigenvalue decomposition of the two-point spatial covariance, yields orthogonal spatial modes that optimally capture the flow energy. The method follows the Karhunen-Loève framework and is also known as principal component analysis [23,24]; it was introduced to fluid mechanics by Lumley [1] and has since been widely used and extended. Although POD is not restricted to spatial coherence, it was initially used mainly to extract spatial structures [25,26]; To distinguish this most common variant of POD from others, the term space-only POD is also commonly used [6].

Sirovich [2] introduced the method of snapshots, which reduces the cost of computing POD by solving an eigenproblem for the temporal correlation (snapshot) matrix that shares the nonzero eigenvalues with the spatial covariance matrix. Borée [27] proposed extended POD to identify, in a second variable, the patterns most correlated with a reference measurement. Sieber *et al.* [28] applied a time-filtered POD, yielding modes whose spectra are concentrated within a frequency band. Mendez *et al.* [29] combined multiresolution (wavelet) analysis with POD; the resulting multiscale POD isolates structures across scales and is not restricted to statistically stationary processes. The natural extension to this is space-time POD, which identifies an energy-optimal basis for the fully space- and time-resolved signal, without *a priori* scale restrictions [1,30,31]. This, however, requires an eigenvalue decomposition of the space-time covariance and is computationally expensive. Modified versions that introduce temporal conditioning to predict extreme events [32] and generate forecasts [33] have been introduced more recently.

Many relevant flows are statistically stationary, so space-time POD is unnecessary, and energy-optimal spatial modes can be identified as a function of frequency [1,34]. These are obtained by an eigendecomposition of the cross-spectral-density (CSD) matrix, a method first termed spectral proper orthogonal decomposition (SPOD) by Picard and Delville [35]. SPOD identifies an energy-optimal basis at each frequency [1]. Its application relies on large datasets to estimate the CSD across frequencies. In practice, Welch segmentation is used: the record is assumed to be ergodic, split into blocks, Fourier modes are computed per block, and the CSD is estimated from the ensemble of blockwise realizations [36]. SPOD is widely applied [35,37–40], also due to its link to resolvent analysis [6], which provides a physics-based model for the data-driven modes. Variants and extensions of SPOD have also been proposed. Colanera *et al.* [41] introduced a robust POD formulation to mitigate sensitivity to outliers, improving mode convergence on noisy datasets. Blanco *et al.* [42] applied a time shift to align advecting structures and thereby improve SPOD mode convergence. Heidt and Colonius [43] extended SPOD to cyclostationary flows in which the broadband statistics are periodic rather than stationary.

Despite its success in identifying temporally and spatially coherent structures, SPOD with Welch averaging inherits the estimation variance–bias tradeoff of spectral estimation from finite records [6,44]. Increasing frequency resolution (longer segments) improves the accuracy of frequency and power estimates for sharp spectral features, but yields high variance and slow convergence in broadband regions. Increasing the number of blocks (shorter segments) reduces variance and improves convergence, but degrades frequency resolution and biases estimates near sharp features. To improve this tradeoff, Schmidt [45] introduced multitaper SPOD, leveraging Thomson’s [17] multitaper estimator. As noted above, multitapering still entails a bias-variance tradeoff, but it outperforms Welch segmentation for PSD estimation at matched effective resolution [19]. To reduce the significant increase in computational cost associated with multitapering, Schmidt [45] proposed a hybrid SPOD that combines Welch segmentation with multitapering. Building on this, Yeung and Schmidt [44] developed an adaptive SPOD formulation inspired by Riedel and Sidorenko [46], based on sinusoidal tapers. The tapered Fourier-mode estimates are obtained from a single

Fourier transform of the signal, using zero padding to twice the record length, which substantially reduces cost. Following the idea of adapting the estimator with frequency [46,47], the number of tapers is selected by an alignment–convergence criterion: broadband regions use more tapers to ensure convergence, whereas near spectral peaks, where tonal energy yields rapid convergence, fewer tapers are used to keep spectral bias low. In this way, the adaptive SPOD method improves convergence of the spectral estimate, reduces variance in broadband regions, and keeps the bias near tonal components low. However, the approach is iterative and may require many eigenvalue decompositions before convergence, which entails high computational cost.

In this study, we present an algorithm for computing SPOD modes based on ensembles of Fourier modes within a frequency band, referred to as band-ensemble SPOD (bSPOD). The idea is inspired by frequency smoothing, which reduces variance in PSD estimates by averaging a full-record periodogram over neighboring frequencies. For standalone PSD estimation, Welch’s segmentation is more commonly used than frequency smoothing. However, when combined with POD in the band-ensemble strategy, the approach unfolds its full potential. The modifications to the commonly used SPOD algorithm with Welch’s ansatz are small, yet the advantages are substantial: bSPOD reduces spectral leakage and provides a data-driven frequency estimate for each mode, enabling variance reduction in broadband regions while keeping bias near tonal eigenvalues low, without relying on an adaptive procedure. Computational cost is comparable to SPOD with Welch segmentation. Since the smoothing is applied after computing the spectrum, bSPOD allows the effective frequency resolution to be modified without requiring recomputation of the Fourier transform.

In the following, we first introduce standard Welch-based SPOD and then describe bSPOD. The methods are compared on an artificial dataset to highlight the advantages of bSPOD and on experimental PIV measurements of a broadband-tonal cavity flow to demonstrate applicability to real data. Results and features are discussed, and the study concludes with a summary of the main findings.

II. METHODOLOGY

In this section, we outline the Welch-based SPOD algorithm and its band-ensemble adaptation. We focus on the time-discrete formulation and refer the reader to the literature for further details [6,34,36,44].

A key focus of the present study is broadband-tonal flow configurations. SPOD estimates a CSD under stationarity assumptions. A true spectral line corresponds to infinite temporal coherence and is therefore fundamentally singular relative to the broadband component. In finite data, the distinction between a line spectrum, a narrowband stochastic process, and cyclostationary dynamics is not identifiable via the Fourier transform. Accordingly, the present methods should be interpreted operationally as a decomposition for finite broadband-tonal datasets, rather than as evidence that the underlying process admits a strictly stationary ergodic spectral representation

A. Spectral proper orthogonal decomposition

We consider a discrete space-time signal $q(x_l, t_n)$, with $l = 0, \dots, N_x - 1$ spatial degrees of freedom and $n = 0, \dots, N_t - 1$ time samples. Here, we use space and time as coordinates as those are the most common encountered in fluid dynamics. Multiple spatial directions and physical variables can be stacked into the l -index. The time series is divided into N_b blocks of length N_w . To increase the number of blocks, adjacent segments may be allowed to overlap. In the following, we focus on the core algorithm without overlapping and consider overlapping separately in Sec. IV E. Each block is transformed into frequency space using a temporal discrete Fourier transform,

$$\hat{q}^{(b)}(x_l, f_j) = \frac{1}{N_w} \sum_{n=0}^{N_w-1} w(t_n) q^{(b)}(x_l, t_n) \exp(-i 2\pi j n / N_w), \quad j = 0, \dots, N_w - 1, \quad (2)$$

where i denotes the imaginary unit, b the block index, j the frequency index, and $w(t_n)$ a taper (window) used to reduce spectral leakage. The hat symbol denotes a complex Fourier amplitude. The frequency spacing results from the length of the windows, $\Delta f = 1/(N_w \Delta t)$, where Δt denotes the time step between two consecutive time samples. For each frequency bin, the Fourier coefficients from the different blocks can be arranged into the data matrix

$$\mathbf{Q}_j = [\hat{\mathbf{q}}_j^{(1)}, \hat{\mathbf{q}}_j^{(2)}, \dots, \hat{\mathbf{q}}_j^{(N_b)}], \quad (3)$$

where each column $\hat{\mathbf{q}}_j^{(b)}$ is a spatially resolved Fourier coefficient at frequency j associated with block b . The shape of $\mathbf{Q}_j$ is $N_x \times N_b$. By averaging over all blocks an estimate of the CSD matrix at frequency j is obtained

$$\mathbf{C}_j = \frac{\alpha}{N_b} \mathbf{Q}_j \mathbf{Q}_j^*, \quad (4)$$

where $(\cdot)^*$ denotes the conjugate transpose. The prefactor is

$$\alpha = \frac{1}{\Delta f U}, \quad U = \frac{1}{N_w} \sum_{n=0}^{N_w-1} w^2(t_n), \quad (5)$$

which normalizes the window energy and ensures that the SPOD eigenvalues have units of power spectral density (variance per unit frequency). The definition of α is consistent with the forward transform normalization in Eq. (2) and ensures that Parseval's theorem is respected [44].

An eigenvalue decomposition of $\mathbf{C}_j$ yields the SPOD modes

$$\mathbf{C}_j \mathbf{W} \Phi_j = \Lambda_j \Phi_j, \quad (6)$$

where $\mathbf{W}$ is a weight matrix that defines the discretized inner product. The eigenvectors contained in the columns of Φ are the SPOD modes that are orthonormal with respect to the inner product, $\Phi^* \mathbf{W} \Phi = \mathbf{I}$. The spectral power density of the modes is contained in the diagonal matrix Λ . If the number of spatial degrees of freedom is larger than the number of blocks, then one typically solves an alternative eigenvalue problem with the same nonzero eigenvalues known as the method of snapshots [2]

$$\frac{\alpha}{N_b} \mathbf{Q}_j^* \mathbf{W} \mathbf{Q}_j \Theta_j = \Lambda_j \Theta_j, \quad (7)$$

where Θ_j are the expansion coefficients that in space-only POD are the time coefficients. This eigenvalue problem is of size $N_b \times N_b$. The SPOD modes follow from

$$\Phi_j = \sqrt{\frac{\alpha}{N_b}} \mathbf{Q}_j \Theta_j \Lambda_j^{-\frac{1}{2}}. \quad (8)$$

This expression shows that the expansion coefficients Θ_j in SPOD quantify how much each block-Fourier coefficient contained in $\mathbf{Q}_j$ contributes to a given SPOD mode in Φ_j .

Overall this yields N_b SPOD modes that can be ranked according to the power density contained in the eigenvalue. All modes are associated with the respective Fourier-bin frequency $f_j = j \Delta f$, with

$$j = \begin{cases} -N_w/2, \dots, +N_w/2 - 1, & N_w \text{ even,} \\ -(N_w - 1)/2, \dots, +(N_w - 1)/2, & N_w \text{ odd.} \end{cases} \quad (9)$$

Compared to the definition in Eq. (2), the frequency index is shifted such that zero frequency is centered. Since the window length N_w is determined by the total duration of the time series N_t and the number of blocks N_b , SPOD faces a tradeoff between frequency resolution and statistical convergence. Increasing the frequency resolution yields fewer Fourier coefficient samples, and therefore less well-converged eigenvalues, which manifests as higher variance in the estimated

modal power spectral density. Conversely, increasing the number of samples by using shorter windows improves statistical convergence but reduces frequency resolution and increases spectral leakage. This results in a loss of spectral sharpness, commonly referred to as spectral bias, which is particularly pronounced when distinct tonal components are smeared into neighboring frequency bins due to leakage. This tradeoff between spectral variance and spectral bias is a well known problem in spectral estimation.

When considering this tradeoff it is important to recall that the SPOD modes form an energy-optimal basis for the dynamics of each frequency bin. Owing to the finite segment length, each bin receives contributions from a range of frequencies through spectral leakage, with the strongest contributions originating from neighboring, unresolved spectral content. One approach to quantify the effective width of this most contributing frequency band is the equivalent noise bandwidth, which for a rectangular window is equal to Δf [48]. Consequently, the additional samples obtained by considering shorter segments only improve the convergence of the spectral estimate if the flow dynamics contributing within the widened effective bandwidth Δf remain sufficiently similar. This requires that the flow dynamics do not vary abruptly with frequency.

We note that tapering decreases spectral leakage by reducing the amplitude of side lobes of the spectral window function at the cost of increasing the main lobe width, which in turn increases the effective bandwidth [48]. In simple terms, tapering reduces long-range leakage but increases local spectral smearing.

B. Band-ensemble spectral proper orthogonal decomposition

The band-ensemble SPOD builds on the frequency-smoothing approach used in spectral estimation that is extended to time–space data. In contrast to the Welch-based algorithm, bSPOD starts with a decisive difference: the data are not segmented into temporal blocks, but the discrete Fourier transform is computed over the entire time series,

$$\tilde{q}(x_l, f_\gamma) = \frac{1}{N_t} \sum_{n=0}^{N_t-1} q(x_l, t_n) \exp(-i 2\pi \gamma n / N_t), \quad \gamma = 0, \dots, N_t - 1. \quad (10)$$

To distinguish symbols from the Welch-based approach, we use a tilde. The frequency spacing of the Fourier transform, $\tilde{\Delta f} = 1/(N_t \Delta t)$, is thus increased by the ratio of total to block length, i.e., N_t/N_w , compared to the Welch-based SPOD. No tapering is applied in the present formulation. While tapering is not fundamentally incompatible with the method, bSPOD relies on a single Fourier transform of the full-length record, for which the increased record length yields a narrower spectral window associated with the implicitly assumed rectangular window and thus naturally reduces spectral leakage compared to the shorter segments used in Welch-based SPOD. Tapering can further suppress the side lobes of the spectral window, but at the cost of broadening its main lobe and thereby increasing correlations between neighboring Fourier coefficients, as already mentioned above.

In contrast to the Welch approach, the data matrix is not constructed from repeated realizations at a fixed frequency, but from N_f consecutive Fourier modes,

$$\tilde{\mathbf{Q}}_j = [\tilde{\mathbf{q}}_{j+1}, \tilde{\mathbf{q}}_{j+2}, \dots, \tilde{\mathbf{q}}_{j+N_f}], \quad j = bN_f, \quad b = 0, \dots, N_t/N_f, \quad (11)$$

where $\tilde{\mathbf{q}}_\gamma$ denotes the full length Fourier mode at frequency γ . To illustrate the connection with the Welch-based formulation, we consider the case where the number of blocks in the Welch formulation equals the number of consecutive Fourier modes considered in the band-ensemble ansatz, $N_b = N_f$, and no overlap is used. In this situation, the Welch-SPOD matrix $\mathbf{Q}_j$ contains N_b realizations of a Fourier mode defined on a frequency grid with spacing Δf , while the band-ensemble matrix $\tilde{\mathbf{Q}}_j$ contains N_f consecutive Fourier modes at an enhanced resolution of $\tilde{\Delta f} = \Delta f / N_f$. The

corresponding analog of the cross-spectral density matrix can be approximated as

$$\tilde{\mathbf{C}}_j \approx \frac{\tilde{\alpha}}{N_f} \tilde{\mathbf{Q}}_j \tilde{\mathbf{Q}}_j^*, \quad (12)$$

where $\tilde{\alpha} = 1/\widetilde{\Delta f}$ is adopted to the increased frequency resolution in the Fourier transform step. With this definition of the prefactors, choosing $N_f = N_b$ implies that both methods consider the dynamics with the same frequency spacing $\Delta f = N_f \widetilde{\Delta f}$, and use the same number of samples per estimate. This can be seen by comparing Eqs. (4) and (12): Without overlap and for $N_f = N_b$, the CSD estimate in Welch-SPOD can be interpreted as the average over N_b realizations, each defined with an effective bandwidth of Δf [see Eq. (4)]; the prefactor becomes $\Delta f/N_b$ assuming no tapering is applied ($U = 1$). In bSPOD, under the same conditions, the prefactor $\tilde{\alpha}/N_f$ in the CSD definition [Eq. (12)] reduces to Δf such that the resulting matrix represents the cross-spectral density with the same effective frequency bandwidth. The remaining factor $1/N_b$ differs between the two methods because bSPOD does not average N_b independent realizations, each defined with frequency spacing Δf , as is done in SPOD with Welch segmentation. Instead, it aggregates the contributions of $N_f (= N_b)$ consecutive Fourier modes within the band Δf , each defined on an N_f -times finer resolution $\widetilde{\Delta f} = \Delta f/N_f$. Consequently, the prefactor ensures that both CSD matrices represent the same total power within the effective bandwidth Δf . At the same time, estimating the CSD from consecutive Fourier modes at higher frequency resolution is advantageous because in-band information can be recovered through the expansion coefficients, as shown in the following.

Due to the importance of the expansion coefficients, bSPOD relies on applying the method of snapshots

$$\frac{\tilde{\alpha}}{N_f} \tilde{\mathbf{Q}}_j^* \mathbf{W} \tilde{\mathbf{Q}}_j \tilde{\mathbf{\Theta}}_j = \tilde{\mathbf{\Lambda}}_j \tilde{\mathbf{\Theta}}_j, \quad (13)$$

which is of size $N_f \times N_f$. The bSPOD modes then follow as

$$\tilde{\mathbf{\Phi}}_j = \sqrt{\frac{\tilde{\alpha}}{N_f}} \tilde{\mathbf{Q}}_j \tilde{\mathbf{\Theta}}_j \tilde{\mathbf{\Lambda}}_j^{-1/2}, \quad (14)$$

so that $\tilde{\mathbf{\Phi}}_j^* \mathbf{W} \tilde{\mathbf{\Phi}}_j = \mathbf{I}$. Analogous to the method of snapshot in Welch-based SPOD, the expansion coefficient in Eq. (14) quantify how much each element of $\tilde{\mathbf{Q}}_j$ contributes to each bSPOD mode contained in $\tilde{\mathbf{\Phi}}_j$. Because in bSPOD each element in $\tilde{\mathbf{Q}}_j$ is a Fourier mode defined on a narrow and unique frequency bin, this information can be exploited to obtain in-band frequency information for each mode. Specifically, we assign each bSPOD mode a representative, data-driven frequency by using the expansion coefficients as frequency weights. The columns of $\tilde{\mathbf{\Theta}}_j$ contain the contribution of each discrete Fourier mode to a given bSPOD mode; we therefore define frequency weights as

$$\beta_{\eta, \gamma} = \frac{|\theta_{\gamma, \eta}|^\ell}{\sum_{\gamma'=1}^{N_f} |\theta_{\gamma', \eta}|^\ell}, \quad \sum_{\gamma'=1}^{N_f} \beta_{\eta, \gamma'} = 1, \quad (15)$$

where $\theta_{\gamma, \eta}$ are the elements of $\tilde{\mathbf{\Theta}}_j$ and ℓ is an exponent to control the type of weighting. The data-driven frequency estimates for the N_f ranked bSPOD modes η are then obtained by a matrix vector product between the weights $\boldsymbol{\beta}$ and the frequencies $\mathbf{f}_j$ of the discrete Fourier modes contained in $\tilde{\mathbf{Q}}_j$ [see Eq. (11)],

$$\tilde{\mathbf{f}}_j = \boldsymbol{\beta} \mathbf{f}_j \quad \mathbf{f}_j = [f_{j+1}, f_{j+2}, \dots, f_{j+N_f}]^\top. \quad (16)$$

Choosing $\ell = 0$, the weights become $1/N_f$ and each mode is attributed to the window center frequency. With $\ell = 1$, the average is weighted by the relative contributions of the Fourier modes to the bSPOD mode. For $\ell = \infty$, each bSPOD mode is assigned the frequency of its single most contributing Fourier mode.

Overall, the data-driven frequency attribution allows each bSPOD mode, though defined over a frequency band, to be associated with a more precise in-band frequency estimate. This is made possible by forming the CSD matrix from neighboring Fourier modes obtained from a single Fourier transform of the full-time record. The key are the expansion coefficients that allow in-band frequency information of the high-resolution Fourier modes to be maintained, when the optimal basis is computed through POD of consecutive Fourier modes.

We emphasize that both Welch-based SPOD and bSPOD seek to approximate the same underlying CSD operator. Consequently, provided that the CSD approximation is sufficiently converged and the underlying assumptions hold, the resulting modes retain the theoretical properties associated with SPOD, including the established connection to resolvent analysis [6]. An open-source MATLAB implementation of bSPOD is available online [49].

C. Band-ensemble filter length

The main parameter that affects bSPOD is the frequency-window length N_f , which, as introduced above, plays a role analogous to the number of blocks, N_b , in Welch-SPOD. For small N_f , the bSPOD modes are computed from only a few consecutive Fourier modes, which improves frequency resolution but reduces statistical convergence, as the eigenvalue problem is small and the dynamics in the considered band must be represented by few modes. Increasing N_f adds consecutive Fourier modes: the modes in total have to represent the dynamics over a wider band, but the larger problem also provides more flexibility to distribute power among the modes. If the frequency content changes slowly across frequency, convergence improves as more Fourier modes are included. Particularly interesting are the limits $N_f = 1$ and $N_f = N_t$. In the limiting case $N_f = 1$, the data matrix $\mathbf{Q}_j$ reduces to a single column,

$$\mathbf{Q}_j = [\tilde{\mathbf{q}}_j], \quad (17)$$

so the corresponding eigenvalue problem becomes trivial and the bSPOD mode coincides with the normalized discrete Fourier mode itself, with its squared amplitude $\tilde{\mathbf{q}}_j^* \tilde{\mathbf{q}}_j$ as the associated eigenvalue. In the opposite limit $N_f = N_t$, all Fourier modes are included and the modes simplify to space-only POD modes. This can be shown by defining the mean-subtracted space-time data matrix

$$\mathbf{X} = [\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{N_t-1}] \in \mathbb{C}^{N_x \times N_t}, \quad (18)$$

with N_x spatial and N_t temporal degrees of freedom, and the unitary DFT matrix

$$\mathbf{F} \in \mathbb{C}^{N_t \times N_t}, \quad F_{l,n} = \frac{1}{\sqrt{N_t}} e^{-i 2\pi n l / N_t}, \quad \mathbf{F}^* \mathbf{F} = \mathbf{I}. \quad (19)$$

This allows us to write the bSPOD data matrix as $\tilde{\mathbf{Q}}_j = \mathbf{X} \mathbf{F}$. Hence the bSPOD cross-spectral density estimate, according to Eq. (12) becomes

$$\tilde{\mathbf{C}}_j = \frac{\Delta t}{N_t} \tilde{\mathbf{Q}}_j \tilde{\mathbf{Q}}_j^* = \frac{\Delta t}{N_t} \mathbf{X} \mathbf{F} \mathbf{F}^* \mathbf{X}^* = \frac{\Delta t}{N_t} \mathbf{X} \mathbf{X}^*, \quad (20)$$

where the prefactor varies slightly from the definition in Eq. (12) since the Fourier transform is based on a unitary normalization. After normalizing with the time step Δt , Eq. (20) becomes the time-averaged covariance matrix. Thus the bSPOD eigenproblem reduces to the space-only POD covariance problem. From this we conclude that the parameter N_f spans the transition from $N_f = 1$, where the modes reduce to DFT modes and capture only temporal coherence, to $N_f = N_t$, where the modes become space-only POD modes that optimally represent the spatial second-order statistics of the entire time series. Intermediate values of N_f yield bSPOD modes that optimally represent the spatial second-order statistics within a local frequency band, and thus structures of temporal and spatial coherence. The same limits are reached for Welch-based SPOD: with a single block (no segmentation) the CSD estimate is rank-1 and the mode equals the normalized DFT mode; with N_t blocks of length one, the CSD matrix becomes the covariance matrix, since the DFT of a single




FIG. 1. Schematic comparison between Welch-based and Band-ensemble SPOD.

time sample is the sample itself, and the SPOD modes become POD modes. Thus, in these parameter limits, bSPOD converges to the same endpoints as other SPOD algorithms, including Welch-based SPOD and the time-filtered POD of Sieber *et al.* [28].

A key advantage of bSPOD is that the band-ensemble filter length N_f does not have to be constant across the spectrum, but may vary with frequency. In practice, shorter bands can be chosen in regions where modes are expected to undergo rapid changes in spatial wavelength, while larger bands can be used where the modes vary more slowly. This is especially relevant at low frequencies, where relative frequency changes are naturally larger. Since spatial wavelengths are often coupled to temporal frequency through a dispersion relation, this can lead to substantial variation in the spatial extent of coherent structures. For turbulent jets, for example, Welch-based SPOD can suffer from this effect, and variable frequency-bin widths have been recommended [50]. The advantage of bSPOD is that the band-ensemble filter length can be adapted locally in the spectrum without recomputing the Fourier transform, allowing the effective bSPOD frequency resolution to vary with frequency. In contrast, the Welch-based SPOD algorithm requires recomputing the Fourier coefficients whenever the frequency resolution (block length) is changed, so modifying the frequency resolution entails complete reprocessing of the data.

D. Summary of differences

As a final part of the methodology section, we summarize how Welch-based SPOD and the bSPOD approach differ. A schematic of how the data matrices are formed in both cases is shown in Fig. 1.

The Welch-based algorithm (left) begins by segmenting the time series into blocks and transforming each block into Fourier space via a discrete Fourier transform (black curved arrows labeled DFT). This yields N_b Fourier mode realizations for each frequency. The Fourier modes at a given frequency are then assembled into data matrices $\mathbf{Q}_j$, from which the CSD matrix is estimated and the SPOD modes are computed [see Eqs. (4)–(8)]. The number of frequency bins N_w results from the total signal length N_t and the number of blocks N_b . The CSD estimates and resulting SPOD

modes are defined at the respective Fourier-bin frequencies and have an equivalent noise bandwidth of Δf , assuming rectangular windows.

bSPOD (right) reverses the order: a single Fourier transform is computed on the full record, yielding Fourier modes on a finely spaced frequency grid ($\widetilde{\Delta f}$). Compared to the Welch-based approach without overlap, the spacing is N_b times smaller. Block matrices $\mathbf{Q}_j$ are then formed from N_f neighboring Fourier modes. The CSD matrix is approximated by summing the contributions of these narrowly spaced modes, and the eigenvalue problem is solved to obtain the bSPOD modes [see Eqs. (12)–(14)]. Similar to the Welch-based method, the number of bSPOD frequency bins $\widetilde{N}_w$ results from the total time length N_t and the filter lengths N_f . The effective bandwidth of the CSD estimate and the corresponding SPOD modes spanned by the ensemble of N_f consecutive Fourier modes is $N_f \widetilde{\Delta f}$.

If the number of consecutive Fourier modes N_f used in bSPOD is equal to the number of blocks N_b in Welch-based SPOD, the resulting CSD matrices have the same size and represent the cross-spectral density with the same effective frequency bandwidth, $\Delta f = N_f \widetilde{\Delta f}$. Under the usual SPOD assumption that the flow dynamics associated with the contributing spectral content remain sufficiently similar, the two approximations should yield the same results.

The key distinction is that Welch-based SPOD relies on temporal segmentation to estimate the CSD matrix, whereas bSPOD estimates modes from contributions of consecutive Fourier coefficients defined on a finer frequency grid obtained from a single Fourier transform of the full-length record. This difference gives rise to several theoretical and practical implications. First, the different approaches used to estimate the CSD lead to a different interpretation of the SPOD expansion coefficients. In Welch-SPOD, the coefficients quantify the contribution of individual time blocks to a mode. In bSPOD, the coefficients instead quantify the contribution of neighboring Fourier coefficients within a frequency band. This frequency-resolved representation enables the frequency-attribution procedure introduced above and provides additional information about the spectral distribution of modal power within a band.

Second, Welch-based SPOD employs block segmentation and assumes statistical independence between blocks. In contrast, bSPOD does not introduce this assumption and instead relies on a Fourier transform of the complete time signal. Consequently, bSPOD relies on a single long observation period over which the statistics remain sufficiently stationary and is not directly suited for combining multiple shorter records of the same process. In the limit of infinite observation time, Fourier coefficients at distinct frequencies are theoretically uncorrelated, thereby providing independent contributions to the approximation of the CSD matrix. While finite record lengths introduce correlations between neighboring Fourier coefficients through spectral leakage, these effects decrease with increasing record length. Consequently, neighboring Fourier coefficients provide approximately independent realizations for sufficiently long time series.

For $N_f = N_b$, both methods define CSD matrices associated with the same effective bandwidth Δf . In a direct comparison, however, it should be noted that Welch-SPOD associates the mode estimates with the Fourier-bin frequencies, Eq. (9), whereas bSPOD defines each estimate over an interval spanned by N_f consecutive Fourier modes. Without overlap and starting the ensembles from the first Fourier mode, this results in the bSPOD frequency grid being shifted by approximately $\Delta f/2$ relative to the Welch-SPOD grid, for which the effective bandwidth is represented only implicitly. In practice, however, this distinction is of little importance, as the objective of both methods is to estimate the modal power spectral density as a function of frequency.

Moreover, the bSPOD algorithm design entails a slightly different computational cost profile. A single Fourier transform of length N_t is more expensive than N_b Fourier transforms of length N_t/N_b . However, bSPOD avoids overhead costs associated with segmentation, tapering, and looping over blocks. Additional computational cost may arise because bSPOD relies on the method of snapshots, which is more expensive when the number of spatial degrees of freedom is small compared to the frequency-window length N_f , and because of the additional operations required for frequency attribution. Nevertheless, the overall computational cost remains comparable to that of Welch-based SPOD, since the eigenvalue problem, typically the most expensive step, is of identical size for

comparable configurations ($N_b = N_f$). In return, bSPOD offers clear advantages: performing the Fourier transform over the full-time series reduces spectral leakage, and the frequency-attribution step enables improved frequency estimations for tonal components as demonstrated in Sec. IV. Last, the effective bSPOD frequency resolution can be varied without recomputing the Fourier transform, making the approach suitable for adaptive frequency resolution strategies [44].

III. DATABASES

A. Artificial broadband-tonal signal

To assess the performance of bSPOD, we construct an artificial space-time signal that is statistically stationary and contains both broadband and tonal features. The goal is to mimic typical characteristics of turbulent flows while retaining full control over the spectral content. The signal is defined on a periodic one-dimensional domain of length L_x , sampled at N_x equispaced points. Time is discretized into N_t samples with time step Δt . The corresponding grids are

$$x_l = l \Delta x, \quad l = 0, \dots, N_x - 1, \quad t_n = n \Delta t, \quad n = 0, \dots, N_t - 1, \quad (21)$$

with $\Delta x = L_x/N_x$ and $\Delta t = T/N_t$. The broadband part of the signal is obtained by superimposing two independent signals y_m , $m = 1, 2$, each representing a convecting wave packet with random phase and amplitude. The target temporal spectrum for both signals is chosen as the response of a single oscillator with natural frequency f_0 and damping ζ :

$$S(\omega) = \frac{4\zeta^2 \omega_0^4}{[\omega_0^2 - \omega^2]^2 + [2\zeta \omega_0 \omega]^2}, \quad (22)$$

where $\omega = 2\pi f$ denotes the angular frequency.

The signal is constructed in frequency space and then inverse transformed to the physical domain. At each discrete frequency $\omega_j = j 2\pi/T$ we draw independent complex Gaussian coefficients,

$$A(\omega_j) = \sqrt{S(\omega_j) \Delta f} z_j, \quad z_j = \frac{(z_{\text{re}} + i z_{\text{im}})}{\sqrt{2}}, \quad (23)$$

where $z_{\text{re}}, z_{\text{im}} \sim \mathcal{N}(0, 1)$ are independent standard normal draws and $\Delta f = 1/T$ is the frequency bin width. This construction ensures $\mathbb{E}[|A(\omega_j)|^2] = S(\omega_j) \Delta f$, so that the expected power spectrum matches the target, while the random phase renders the realization statistically stationary. The scaling with Δf ensures that the variance of the signal equals the integral of the prescribed power spectral density. Consequently, apart from the random realization, both broadband signals share the same spectral distribution, while differences arise primarily from their spatial structure, which is considered next.

To introduce spatial coherence, each frequency component is assigned a wave number through the linear dispersion relation

$$k_m(\omega_j) = \frac{\omega_j}{c_m} + k_{0,m}, \quad (24)$$

modeling convection at speed c_m with offset $k_{0,m}$, which are chosen differently for the two broadband components. To avoid introducing a preferred spatial origin, the two broadband branches are displaced by a random shift $x_{0,m} \sim \mathcal{U}(0, L_x)$, where $\mathcal{U}$ denotes the uniform distribution. This ensures that the ensemble is statistically homogeneous in space. The resulting signals read

$$y_m(x_l, t_n) = \sum_{j=1}^{N_t/2} A_m(\omega_j) \exp(i[k_m(\omega_j)(x_l - x_{0,m}) + \omega_j t_n]). \quad (25)$$

To complement the broadband signal, we superimpose five discrete tones:

$$o_r(x_l, t_n) = B_r \exp(i[k_r x_l / L_x + \omega_r t_n + \phi_r]), \quad r = 1, \dots, 5, \quad (26)$$




FIG. 2. Mach 0.6 cavity-flow mono-PIV dataset of Zhang *et al.* [51]. Left: time-averaged axial (top) and vertical (bottom) velocity field; thick black lines mark the cavity boundaries. Right: snapshot of the fluctuating (mean-subtracted) vertical velocity (top) and the PSD (bottom) of the fluctuating vertical velocity at the green dot. Red vertical lines indicate the three Rossiter modes (940, 1540, and 2130 Hz).

with amplitude B_r , frequency ω_r , spatial wave number k_r , and random phase $\phi_r \sim \mathcal{U}(0, 2\pi)$. These time series generate sharp spectral peaks adding tonal components to the broadband signals. The complete artificial signal is obtained by superimposing the broadband and the tonal components

$$y(x_l, t_n) = \sum_{m=1}^2 g_m y_m(x_l, t_n) + \sum_{r=1}^5 o_r(x_l, t_n), \quad (27)$$

with prescribed scalar weights g_m for the two broadband signals to separate them into a leading and a sub-leading mode.

In the results shown, we set $N_x = 250$, $L_x = 10$, $N_t = 100\,000$, $\Delta t = 10^{-3}$. For the two broadband signals, we use $f_0 = 100$, $\zeta = 0.2$, $(g_m, c_m, k_{0,m}) \in \{(1.0, 200, 0.5), (0.5, 150, 1.5)\}$. Five tones are added with $(B_r, f_r, k_r) \in \{(5, 23, 3), (3, 25, 8), (6, 30, 5), (5, 75, 3), (5, 125, 6)\}$.

B. Broadband-tonal cavity flow

To demonstrate the methodology on experimental data, bSPOD is applied to a high-speed mono-PIV data set of a turbulent open cavity flow [51]. This data set has previously been used to benchmark spectral estimation techniques, for example the adaptive SPOD algorithm based on multitaper estimation for improved tone identification [44], and by Zhang *et al.* [51], who estimate spectral POD modes of non-time-resolved PIV snapshots by making use of high-speed pressure measurements.

The cavity has a depth of $H = 26.5$ mm, an axial extent of $6H$, and an out-of-plane width of $3.85H$. The PIV measurements provide the center-plane axial and vertical velocity components on an equidistant grid of $N_x = 154$ by $N_y = 54$ points, with a total of $N_t = 16\,000$ snapshots sampled at $\Delta t = 6.25 \times 10^{-5}$ s. Figure 2 shows the time-averaged axial and vertical velocity fields, a snapshot of the fluctuating vertical velocity component u'_y , and the spectral power density of the fluctuating vertical velocity component $u'_y(2, 0)$. The location at which this component is extracted is highlighted with a marker in the u'_y plot on the top right. The PSD is computed using Welch's method with a window length of 1600 samples and 50% overlap. Three distinct peaks are observed at 940, 1540, and 2130 Hz, corresponding to the second, third, and fourth Rossiter modes.

IV. RESULTS

A. Characteristics of bSPOD

To highlight the features of bSPOD, we first consider the artificial broadband-tonal data set. The top panel in Fig. 3 shows the eigenvalue spectra of the first three modes from both Welch-based SPOD and bSPOD, together with the analytic ground-truth values known from the construction of the signal. For the broadband components y_m , the spatially integrated power spectral density is shown; for the tonal components o_r , the integrated tonal power is shown, since their energy is concentrated at single frequencies. SPOD and bSPOD results are shown for $N_b = N_f = 50$, which yields a fine frequency resolution of $\Delta f = N_f \widetilde{\Delta f} = 0.5$ Hz. Both methods are initially applied without overlap or windowing to compare the core algorithms without additional features. Within each frequency band, the bSPOD modes are assigned to frequencies according to Eq. (16) with weight exponent $\ell = 1$. Since the tonal component's spectral densities correspond to Dirac impulses, we scale the associated peak eigenvalues from both SPOD variants by the frequency spacing Δf to enable direct comparison with the tones' total power. The corresponding eigenvalues representing tonal power are enclosed by a red box in Fig. 3.

Both methods accurately capture the frequency and power of the tones. The spectra of the two broadband components are also represented well, although a high estimation variance is observed, visible in the noisy distribution of the eigenvalue estimates (markers) around the expected ground-truth values (black solid lines). The high variance results from the small number of blocks and the short band-ensemble filter width $N_b = N_f$, reflecting the expected tradeoff between frequency resolution and statistical convergence.

The leading and first subleading broadband modes are reproduced by the respective leading and subleading SPOD (bSPOD) eigenvalues. In the bins containing tones, the dominant SPOD (bSPOD) mode represents the tone, shifting the mode ranking by one: the leading broadband component appears as the first subleading mode, the subleading broadband component by the second subleading mode, and so on. These shifted SPOD and bSPOD eigenvalues are highlighted in the top panel of Fig. 3 with green marker edges.

Overall, the differences between the shown SPOD and bSPOD results are small, except that bSPOD shows less spectral leakage. This can be observed from the magnitude of the second subleading mode (square markers), which is purely a result of leakage and is about one order of magnitude lower for bSPOD than for SPOD. This reduction follows from using the full-time series in the Fourier transform, as opposed to shorter windows in Welch-based SPOD with time segmentation.

To reduce the estimation variance, the number of blocks and the band-ensemble filter length are increased by a factor of four, which decreases the frequency resolution of both methods by the same factor. The corresponding eigenvalue spectra are shown in the middle plot of Fig. 3. For bSPOD the results are plotted again for $\ell = 1$, Eq. (15). As in the upper figure with higher frequency resolution, bSPOD reproduces both tonal and broadband components accurately, now with significantly reduced variance for the broadband signals. By contrast, the Welch-based SPOD eigenvalues deviate from the analytic reference: the frequency resolution is insufficient to resolve the tones, and the short windows lead to severe leakage effects. This affects not only the regions near the tonal peaks but also the subleading modes at higher frequencies, where leakage obscures the separation between physical modes and spurious ones.

In Welch-based SPOD, spectral leakage is commonly mitigated by applying a taper to each time segment. The blue markers in the bottom panel of Fig. 3 show the corresponding eigenvalues obtained when a Hann window is used as the taper. The bSPOD eigenvalues are unchanged relative to the middle panel and are replotted again for direct comparison. Tapering substantially reduces spectral leakage for frequencies away from the tonal ones; however, the spurious secondary subleading-mode branch induced by leakage remains more pronounced for Welch-based SPOD than for bSPOD. Furthermore, a side effect of tapering in Welch-based SPOD is increased leakage in the immediate vicinity of the tones. This follows because windowing corresponds to convolution with




FIG. 3. Modal power spectral densities and tonal power for the artificial data set. Welch-based SPOD and bSPOD are compared against the theoretical values for three configurations: $\Delta f = 0.5$, Hz with $N_b = N_f = 50$ (top), $\Delta f = 2$ Hz with $N_b = N_f = 200$ (middle), and $\Delta f = 2$ Hz with $N_b = N_f = 200$ using a Hann taper for Welch-based SPOD (bottom). Blue and orange markers denote the leading three eigenvalues from Welch-based SPOD and bSPOD, respectively. Markers enclosed by the red box represent tonal power (right axis), whereas all other markers and lines indicate power spectral densities (left axis). Black lines show the analytic broadband power spectral density ($g_m^2 S(\omega) L_x$), and black squares indicate the true tonal power ($B_r^2 L_x$; see Sec. III A). Dashed vertical lines mark the tonal frequencies.

the window's frequency response: compared to the rectangular window, the Hann window has a broader main lobe, which increases spectral smearing near the tonal peaks [48].

In bSPOD, leakage is naturally reduced by computing the Fourier transform over the full-time series, without the need for tapering. Applying a taper within bSPOD would further reduce the

amplitude of the spurious mode branch. Overall, these results demonstrate that bSPOD yields improved frequency selectivity and reduced spectral leakage relative to the Welch-based formulation, even when tapering is applied.

B. Application to experimental data

After applying the proposed method to an artificial test data set, we turn to the more relevant case and apply Welch-based SPOD and bSPOD to the mean-subtracted cavity flow velocity snapshots, u'_x and u'_y . For a fair comparison, the methods are again configured with $N_f = N_b$; that is, the number of blocks in Welch-based SPOD equals the number of consecutive Fourier modes considered in bSPOD. No overlap is applied, ensuring that both methodologies produce the same number of mode estimates with the same effective frequency spacing Δf . For Welch-based SPOD, a Hann window is used to reduce spectral leakage, while in bSPOD the frequency attribution is based on $\ell = \infty$; see Eq. (15).

The top panel of Fig. 4 shows the power spectral density of the three most energetic modes obtained with both methods for a frequency resolution of $\Delta f = 25$ Hz ($N_b = N_f = 25$). The dashed vertical lines highlight the frequencies of the three Rossiter modes, obtained from a point-wise PSD as described in Sec. III B. At this high frequency resolution, the three tonal Rossiter peaks are captured well by both approaches. However, the high frequency resolution comes at the cost of large estimation variance, visible as elevated noise levels in regions away from the tones where the leading mode represents broadband dynamics. The observed noise level indicates that the spectra are not well converged and more modes are required which, however, would reduce the frequency resolution. This tradeoff is illustrated in the middle and bottom panels of Fig. 4, which show the same spectra for $N_b = N_f = 50$ ($\Delta f = 50$ Hz) and $N_b = N_f = 100$ ($\Delta f = 100$ Hz), respectively. Although both methods experience the same decrease in frequency resolution, bSPOD maintains accurate tonal frequency estimates enabled by the data-driven frequency attribution. In contrast, such in-band frequency information is not available for Welch-SPOD. Furthermore, the Welch approach remains prone to spectral leakage: energy from neighboring tones contaminates the spectral estimate, resulting in a biased representation of the tonal components. These results highlight the advantages of bSPOD for analyzing broadband-tonal flows. By reducing spectral leakage and exploiting the information contained in the expansion coefficients, bSPOD achieves sharper spectral localization and a clearer separation of tonal and broadband contributions than the Welch-based SPOD formulation.

C. Mode convergence

Having discussed the eigenvalue spectra and their convergence behavior, we now focus on the mode shapes. The modes corresponding to eigenvalues marked with yellow arrows and circles in Fig. 4 are shown in Figs. 5 and 6. The plots compare bSPOD modes computed from N_f neighboring Fourier modes with Welch-SPOD modes evaluated at the lower-bound Fourier-bin frequency of the corresponding interval. Comparisons with modes evaluated at the upper frequency bound lead to the same conclusions and are therefore omitted for brevity.

We begin with a tonal component at 1540 Hz for the fine frequency-resolution case, $\Delta f = 25$ Hz, corresponding to the eigenvalue spectrum shown in the top panel of Fig. 4. Figure 5 shows the real part of the vertical-velocity mode, with the bSPOD result on the left and the Welch-SPOD result on the right. As expected, the two mode shapes are nearly identical, since they are estimated from CSD matrices of equal size and are associated with the same frequency range (1525–1550 Hz). Interestingly, both modes appear well converged despite the high variance in the spectra; see top panel in Fig. 4. This is because the tonal component carries most of the power in that frequency bin, producing a sharply defined contribution to the cross-spectral density and, in turn, a clearly identifiable eigenvector, even at high frequency resolution.




FIG. 4. Leading three SPOD eigenvalues versus frequency for the Welch-based (blue) and band-ensemble (orange) formulations with $N_f = N_b$ and three frequency resolutions: $\Delta f = 25$ Hz (top), 50 Hz (middle), and 100 Hz (bottom). Vertical red lines denote the Rossiter-mode frequencies identified from the single-point PSD (see Fig. 2). Yellow arrows and circles mark the eigenvalues whose corresponding mode shapes are shown in Sec. IV C.

A different behavior is observed for broadband dynamics of lower power. The top row of Fig. 6 shows the real part of the vertical-velocity mode for SPOD and bSPOD at the same frequency resolution, in the 425–450 Hz band where the leading mode is broadband (see the top panel of Fig. 4). Here, the high frequency resolution and consequently small number of Fourier modes




FIG. 5. Real part of the vertical velocity component of the Rossiter mode at 1540 Hz obtained with band-ensemble SPOD (left) and Welch-based SPOD (right), both using an effective frequency resolution of $\Delta f = 25$ Hz. The corresponding eigenvalues are marked by yellow crosses in the top panel of Fig. 4.

($N_f = N_b = 25$) contributing to the estimation of the CSD matrix lead to noisy and poorly converged mode estimates. Importantly, the SPOD and bSPOD modes exhibit a similar level of noise and overall quality, indicating that neither method provides a clear advantage in this case. The bottom row of Fig. 6 shows the corresponding mode shapes for a lower frequency resolution, $\Delta f = 100$ Hz, such that the mode represents the leading dynamics in the 400–500 Hz band. For both methods, the increase in frequency spacing increases the size of the eigenvalue problem ($N_f = N_b = 100$), implying that more orthogonal modes contribute to the representation of the dynamics within the band. As commonly observed for SPOD this reduces the noise level seen in Fig. 6 equally for the Welch and the band-ensemble approach.

It is noted, that this improvement is a result of the dominant flow dynamics not varying substantially in the considered frequency bands; for a finite dataset, increasing the number of realizations, either by adding additional blocks (SPOD) or consecutive Fourier modes (bSPOD), necessarily increases the effective frequency range contributing to each spectral estimate. If the mode shape varies strongly with frequency, for example due to strong dispersion of the dominant wavelength, increasing the effective spacing mixes dissimilar dynamics and can offset the benefit of added realizations. In this case improved convergence can only be achieved by considering a longer time series.

To demonstrate that both methods also yield very similar mode shapes for the subleading modes, Fig. 7 shows the real part of the first subleading bSPOD mode in the 800–900 Hz band and compares it with the corresponding SPOD mode computed at 800 Hz. Both modes exhibit some noise, which is expected due to the lower energy of the subleading mode. Nevertheless, their overall mode shapes are very similar.




FIG. 6. Real part of the vertical velocity component of the leading broadband mode computed at two effective frequency resolutions: 425–450 Hz (top) and 400–500 Hz (bottom). Results are shown for band-ensemble SPOD (left) and Welch-based SPOD (right). The corresponding eigenvalues are marked by yellow crosses in the top ($\Delta f = 25$ Hz) and bottom ($\Delta f = 100$ Hz) panels of Fig. 4.




FIG. 7. Real part of the vertical velocity component of the first subleading broadband mode in the 800–900 Hz band. Results are shown for band-ensemble SPOD (left) and Welch-based SPOD (right). The corresponding eigenvalues are marked by yellow crosses in the bottom panel of Fig. 4 ($\Delta f = 100$ Hz).

Finally, Fig. 8 compares the mode convergence behavior of SPOD and bSPOD with respect to the amount of available data. Therefore the dataset was progressively shortened from 100% to 25% in steps of 25%, while keeping the frequency spacing constant at $\Delta f = 100$ Hz. This implies that the number of modes representing the dynamics within each effective frequency band decreases proportionally with data reduction.

The top row in Fig. 8 shows the alignment between the leading modes (largest eigenvalue) obtained from the reduced, ϕ_{red} , and full datasets, ϕ_{ref} , defined as

$$A = \frac{|\phi_{\text{red}}^* \phi_{\text{ref}}|}{\|\phi_{\text{red}}\| \|\phi_{\text{ref}}\|}, \quad (28)$$

which is a commonly used measure of modal similarity [52]. An alignment value of 1 means that the modes are fully aligned, while a value of 0 indicates that the mode shapes are orthogonal. The bottom row shows the respective alignment for the first subleading modes. As expected, Fig. 8 shows that the alignment increases with the amount of data considered for both methods.




FIG. 8. Mode convergence of the leading (top) and first subleading (bottom) SPOD eigenvalues assessed by reducing the number of snapshots to 75, 50, and 25% of the full-time record. Results are shown for band-ensemble SPOD (left) and Welch-based SPOD (right). Convergence is quantified by keeping the frequency resolution fixed and aligning the leading SPOD mode shape computed from reduced time sequences with the corresponding mode obtained from the full-time record [see Eq. (28)].

Although the exact values differ for individual frequencies, SPOD and bSPOD exhibit comparable convergence behavior. The overall alignment at a given level of data reduction is lower for the first subleading mode (bottom) than for the leading mode (top), indicating that convergence of the subleading mode generally requires more data. This is consistent with its lower energy content and, consequently, lower signal-to-noise ratio. Nevertheless, the overall convergence trends of SPOD and bSPOD also remain very similar for the subleading mode.

We conclude that, under equivalent conditions $N_b = N_f$, corresponding to the same effective frequency bandwidth, the dynamics are represented by the same number of POD modes, and SPOD and bSPOD yield similar mode shapes with comparable convergence characteristics.

D. Frequency attribution

So far we have seen that bSPOD achieves mode convergence comparable to Welch-based SPOD while reducing spectral leakage and improving the frequency estimates of the eigenvalues via the frequency-attribution step. This advantage arises because bSPOD forms the CSD estimate from consecutive Fourier modes, so the method-of-snapshots eigenvalues [the expansion coefficients Θ , Eq. (13)] quantify how much each discrete Fourier mode contributes to a given bSPOD mode [Eq. (14)]. After normalization, the coefficients act as frequency weights which bSPOD uses to infer an in-band frequency estimate for each mode [Eq. (16)]. In this section, we elaborate on the frequency attribution step for the broadband cavity flow.

Figure 9 shows the frequency weights β of the leading cavity-flow mode for $\ell = 1$ and two resolutions, $\Delta f = 50$ and $\Delta f = 100$ Hz (top and bottom). The corresponding eigenvalue spectra are shown above in the middle and bottom panels of Fig. 4. Weights are plotted against the discrete Fourier frequencies contributing to each band between 800 and 1700 Hz; even and odd bands are shown in black and gray to emphasize the band-wise definition. In most bands, the weights fluctuate around the mean $1/N_f$ (green horizontal line) without a distinct pattern. In bands containing a tone, marked by red vertical lines, however, weights near the tone are elevated and contribute more strongly to the bSPOD mode, reflecting the tone's higher power and clear separability from the surrounding broadband dynamics.

The blue dashed lines mark bSPOD frequency estimates from the attribution step using $\ell = \infty$ [Eq. (16)], selecting the frequency from the single most contributing Fourier mode per band. In tonal bands these lie very close to the Rossiter frequencies (red). At $\Delta f = 50$ Hz, both shown Rossiter frequencies lie near an upper band limit. The algorithm accounts for this and assigns frequencies (blue) close to the tone (red) not only to the band containing the tone but also to the adjacent band by placing an attributed frequency at the lower edge of the neighboring band. This yields spikier modal power densities (Fig. 6, middle) than those obtained with Welch-based SPOD, which uses a fixed frequency grid.

The yellow dashed lines in Fig. 9 show the bSPOD frequency estimates obtained with the shown weights ($\ell = 1$). These are less accurate at representing the tonal frequencies than the $\ell = \infty$ estimates, but they still provide better frequency estimates than a fixed choice such as the band-center frequency. In broadband bands without tones, the $\ell = 1$ estimates tend to select frequencies closer to the band center, consistent with power being distributed across the band. In these broadband sections, $\ell = \infty$ can yield frequency estimates near the band edges even when the contributions are nearly uniform. Hence, it is desirable to rely on $\ell = \infty$ for tonal bands and $\ell = 1$ for broadband bands; such a switch could be triggered by the local variation of the eigenvalue power spectral density with frequency. In all cases, the normalization of the frequency weights, Eq. (15), ensures that the estimated frequency remains bounded within the respective band, so neither $\ell = 1$ nor $\ell = \infty$ yields an out-of-band frequency estimate. Consequently, in contrast to Welch-based SPOD, which relies on a fixed frequency grid, the frequency estimation can only improve.




FIG. 9. Frequency weights β of the leading band-ensemble mode versus frequency for $\ell = 1$ and $\Delta f = 50$ Hz (top) and 100 Hz (bottom); see Eq. (15). The alternating gray–white background and alternating black/gray curves indicate the frequency bands used in bSPOD. Vertical blue and yellow lines mark the bSPOD frequency estimates obtained with $\ell = 1$ and $\ell = \infty$, respectively. Red vertical lines denote the Rossiter-mode frequencies (see Fig. 2).

E. Overlap in band-ensemble SPOD

For Welch-based SPOD, it is common to overlap the time segments to increase the number of Fourier estimates from a finite time signal. The same principle can be used in bSPOD by shifting the band-ensemble window by less than N_f , so that consecutive frequency-blocks overlap. This produces additional estimates and thereby refines the frequency grid of the modal power spectrum.

When applying overlap, however, two points must be considered. First, although the bSPOD modes are defined with the frequency spacing associated with the chosen band length N_f , the effective spacing for the power estimate changes with overlapping because Fourier modes are reused across overlapping blocks. In practice, the overlap factor resulting in increased frequency resolution must be considered when integrating the power density over frequency to recover the full power of the signal.

Second, in broadband signals containing tonal components, overlap can artificially spread the energy of a tone across several bSPOD bins if all modes are simply assigned to the band-center frequency ($\ell = 0$). For PSD estimates using frequency smoothing, this produces block-shaped jumps in the spectrum, likely a reason why PSDs are more commonly estimated using Welch segmentation rather than frequency smoothing. In bSPOD, this issue is avoided by the data-driven frequency attribution step, enabled by the POD expansion coefficients. It ensures that each bSPOD mode is linked to its most contributing frequencies and leads to the tonal power being correctly concentrated.




FIG. 10. Leading three bSPOD eigenvalues versus frequency for the cavity-flow data, computed with $N_f = 100$ ($\Delta f = 100$ Hz) and 75% band overlap. Eigenvalues are shown for a fixed frequency grid ($\ell = 0$, assigning each eigenvalue to the band center) and for two data-driven frequency attributions, $\ell = 1$ and $\ell = \infty$ [see Eq. (15)].

To illustrate this behavior, Fig. 10 shows the leading three bSPOD eigenvalues of the cavity-flow data for $N_f = 100$ ($\Delta f = 100$ Hz; compare with the bottom panel of Fig. 4) and a band overlap of 75%. Compared to the case without overlap, the frequency resolution is increased by a factor of four. It should be noted that adjacent modes partially share the same Fourier components, leading to statistically correlated estimates. Results are shown for frequency attribution using $\ell = \infty$ and $\ell = 1$, as well as for the case without attribution, where each bSPOD mode is assigned the band-center frequency ($\ell = 0$).

As expected, without frequency attribution ($\ell = 0$) the spectrum resembles classical frequency smoothing: tonal peaks appear as block-wise elevations. With $\ell = \infty$, the tonal modes are attributed frequencies close to the Rossiter frequencies, thereby preserving sharp tonal peaks in the spectrum despite the band overlap. In the broadband parts of the spectrum, however, $\ell = \infty$ can yield uneven or discontinuous frequency trends, including small jumps, which are particularly apparent for the first subleading mode in Fig. 10. Using $\ell = 1$ produces a smooth and well-resolved broadband spectrum, but yields less sharply localized tonal peaks than $\ell = \infty$, although still improved compared to $\ell = 0$. These observations suggest a hybrid strategy in which $\ell = 1$ is used in broadband bands and $\ell = \infty$ in bands containing tones. As mentioned above, such a switch could be triggered, for example, by detecting strong local variations of the modal power spectral density with frequency.

We conclude that, in bSPOD, introducing overlap between neighboring frequency bands can increase the effective frequency resolution of the modal spectrum without smearing tonal components across adjacent bins, provided that an appropriate frequency attribution is used. In Welch-based SPOD, overlap is typically introduced to increase the number of blocks of fixed length. Alternatively, overlap in Welch-based SPOD can be interpreted as allowing for increased block length at a fixed number of blocks, leading to a comparable effect on frequency resolution. The key advantage of bSPOD is that overlap is applied after the Fourier transform and can be varied across the spectrum, whereas SPOD based on time segmentation relies on a fixed frequency grid and does not offer this flexibility.

V. CONCLUSIONS

This work presents a method for computing SPOD modes of statistically stationary data. The proposed approach, termed band-ensemble SPOD, is inspired by frequency smoothing, a method to reduce estimation variance in power spectral density estimates. Here, this methodology is extended to SPOD, enabling the estimation of spatial modes together with their modal power spectral

densities. Unlike the more commonly applied Welch-based SPOD formulation, which estimates the cross-spectral density matrix from Fourier modes defined on the same frequency bin obtained through temporal segmentation, bSPOD approximates the cross-spectral density matrix, and thus the SPOD modes, from narrowly spaced, neighboring Fourier coefficients. These coefficients are obtained from a single Fourier transform of the full-time series, without segmentation.

Compared with Welch-based SPOD, bSPOD offers several advantages. First, spectral leakage is naturally reduced by avoiding time segmentation. Second, the contribution of individual Fourier coefficients to each bSPOD mode can be tracked through the expansion coefficients, allowing in-band frequency information to be retained. This information can be exploited to obtain data-driven frequency estimates for individual bSPOD modes. By combining these properties, bSPOD reduces estimation variance while maintaining low bias for tonal components, providing a favorable balance in the bias-variance tradeoff relative to the Welch-based formulation. As a result, bSPOD is particularly well suited for broadband-tonal flows, where variance reduction without compromising the accuracy of tonal eigenvalue estimates is desirable.

The bSPOD formulation further allows this tradeoff to be adjusted by varying the filter length with frequency. This enables modal decompositions over variable frequency bands with a frequency-dependent number of modes. Such flexibility makes bSPOD well suited for adaptive strategies that select the number of modes based on convergence criteria, as suggested by Yeung and Schmidt [44]. Given these properties and a computational cost comparable to Welch-based SPOD, bSPOD provides a practical improvement for spectral modal analysis of turbulent flows.

ACKNOWLEDGMENT

This work was supported by a postdoc fellowship of the German Academic Exchange Service (DAAD).

DATA AVAILABILITY

The code that support the findings of this article are openly available [49].

-
- [1] J. L. Lumley, The structure of inhomogeneous turbulent flows, in *Atmospheric Turbulence and Radio Wave Propagation* (Nauka Publishing House, Moscow, 1967), pp. 166–178.
 - [2] L. Sirovich, Turbulence and the dynamics of coherent structures. I. Coherent structures, *Q. Appl. Math.* **45**, 561 (1987).
 - [3] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, Spectral analysis of nonlinear flows, *J. Fluid Mech.* **641**, 115 (2009).
 - [4] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
 - [5] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley, Modal analysis of fluid flows: An overview, *AIAA J.* **55**, 4013 (2017).
 - [6] A. Towne, O. T. Schmidt, and T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *J. Fluid Mech.* **847**, 821 (2018).
 - [7] P. Jordan and T. Colonius, Wave packets and turbulent jet noise, *Annu. Rev. Fluid Mech.* **45**, 173 (2013).
 - [8] S. Demange, Z. Yuan, S. Jekosch, A. Hanifi, A. V. G. Cavalieri, E. Sarraji, T. L. Kaiser, and K. Oberleithner, Resolvent model for aeroacoustics of trailing edge noise, *Theor. Comput. Fluid Dyn.* **38**, 163 (2024).

- [9] K. Sarras, C. Tayeh, V. Mons, and O. Marquet, Linear stability analysis of turbulent mean flows based on a data-consistent Reynolds-averaged Navier-Stokes model: Prediction of three-dimensional stall cells around an airfoil, *J. Fluid Mech.* **1001**, A41 (2024).
- [10] T. J. Poinso, A. C. Troune, D. P. Veynante, S. M. Candel, and E. J. Esposito, Vortex-driven acoustically coupled combustion instabilities, *J. Fluid Mech.* **177**, 265 (1987).
- [11] J. G. R. von Saldern, J. P. Beuth, J. M. Reumschüssel, A. Jaeschke, C. O. Paschereit, and K. Oberleithner, Low-frequency streaky structures in turbulent hydrogen jet flames, *Combust. Flame* **278**, 114231 (2025).
- [12] C. H. Williamson and R. N. Govardhan, Vortex-induced vibrations, *Annu. Rev. Fluid Mech.* **36**, 413 (2004).
- [13] N. Wiener, Generalized harmonic analysis, *Acta Math.* **55**, 117 (1930).
- [14] A. Khintchine, Korrelationstheorie der stationären stochastischen Prozesse, *Math. Ann.* **109**, 604 (1934).
- [15] M. S. Bartlett, Periodogram analysis and continuous spectra, *Biometrika* **37**, 1 (1950).
- [16] P. Welch, The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- [17] D. Thomson, Spectrum estimation and harmonic analysis, *Proc. IEEE* **70**, 1055 (1982).
- [18] B. Babadi and E. N. Brown, A review of multitaper spectral analysis, *IEEE Trans. Biomed. Eng.* **61**, 1555 (2014).
- [19] T. Bronez, On the performance advantage of multitaper spectral analysis, *IEEE Trans. Signal Process.* **40**, 2941 (1992).
- [20] P. J. Daniell, Discussion on symposium on autocorrelation in time series, Supplement to the *J. Roy. Statist. Soc. Supp.* **8**, 85 (1946).
- [21] M. S. Bartlett, Smoothing periodograms from time-series with continuous spectra, *Nature (London)* **161**, 686 (1948).
- [22] R. B. Blackman and J. W. Tukey, The measurement of power spectra from the point of view of communications engineering—Part I, *Bell Syst. Tech. J.* **37**, 185 (1958).
- [23] K. Pearson, LIII. On lines and planes of closest fit to systems of points in space, *Lond. Edinb. Dubl. Phil.* **2**, 559 (1901).
- [24] H. Hotelling, Analysis of a complex of statistical variables into principal components, *J. Educ. Psychol.* **24**, 417 (1933).
- [25] G. Berkooz, P. Holmes, and J. L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.* **25**, 539 (1993).
- [26] J. Weiss, A tutorial on the proper orthogonal decomposition, in *Proceedings of the AIAA Aviation Forum* (AIAA, Reston, VA, 2019).
- [27] J. Borée, Extended proper orthogonal decomposition: A tool to analyse correlated events in turbulent flows, *Exp. Fluids* **35**, 188 (2003).
- [28] M. Sieber, C. O. Paschereit, and K. Oberleithner, Spectral proper orthogonal decomposition, *J. Fluid Mech.* **792**, 798 (2016).
- [29] M. A. Mendez, M. Balabane, and J.-M. Buchlin, Multi-scale proper orthogonal decomposition of complex fluid flows, *J. Fluid Mech.* **870**, 988 (2019).
- [30] S. V. Gordeyev and F. O. Thomas, A temporal proper decomposition (TPOD) for closed-loop flow control, *Exp. Fluids* **54**, 1477 (2013).
- [31] P. Frame and A. Towne, Space-time POD and the Hankel matrix, *PLoS ONE* **18**, e0289637 (2023).
- [32] O. T. Schmidt and P. J. Schmid, A conditional space-time pod formalism for intermittent and rare events: Example of acoustic bursts in turbulent jets, *J. Fluid Mech.* **867**, R2 (2019).
- [33] O. T. Schmidt, Data-driven forecasting of high-dimensional transient and stationary processes via space-time projection, *Proc. R. Soc. A* **482**, 20250454 (2026).
- [34] J. L. Lumley, *Stochastic Tools in Turbulence* (Academic Press, San Diego, CA, 1970).
- [35] C. Picard and J. Delville, Pressure velocity coupling in a subsonic round jet, *Int. J. Heat Fluid Flow* **21**, 359 (2000).
- [36] O. T. Schmidt and T. Colonius, Guide to spectral proper orthogonal decomposition, *AIAA J.* **58**, 1023 (2020).

- [37] J. H. Cintriniti and W. K. George, Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition, *J. Fluid Mech.* **418**, 137 (2000).
- [38] C. Chen, S. Wang, and S. Ghaemi, Spectral proper orthogonal decomposition of time-resolved three-dimensional flow measurements in the turbulent wake of the Ahmed body, *J. Fluid Mech.* **985**, A19 (2024).
- [39] J. G. R. von Saldern, O. T. Schmidt, P. Jordan, and K. Oberleithner, On the role of eddy viscosity in resolvent analysis of turbulent jets, *J. Fluid Mech.* **1000**, A51 (2024).
- [40] B. Steinfurth, M. Li, F. Scarano, and J. Weiss, Three-dimensional low-frequency dynamics of a turbulent separation bubble, *J. Fluid Mech.* **1019**, R4 (2025).
- [41] A. Colanera, O. T. Schmidt, and M. Chiatto, Robust spectral proper orthogonal decomposition, *Comput. Phys. Commun.* **307**, 109432 (2025).
- [42] D. C. Blanco, E. Martini, K. Sasaki, and A. V. Cavalieri, Improved convergence of the spectral proper orthogonal decomposition through time shifting, *J. Fluid Mech.* **950**, A9 (2022).
- [43] L. Heidt and T. Colonius, Spectral proper orthogonal decomposition of harmonically forced turbulent flows, *J. Fluid Mech.* **985**, A42 (2024).
- [44] B. C. Y. Yeung and O. T. Schmidt, Adaptive spectral proper orthogonal decomposition of broadband-tonal flows, *Theor. Comput. Fluid Dyn.* **38**, 355 (2024).
- [45] O. T. Schmidt, Spectral proper orthogonal decomposition using multitaper estimates, *Theor. Comput. Fluid Dyn.* **36**, 741 (2022).
- [46] K. Riedel and A. Sidorenko, Minimum bias multiple taper spectral estimation, *IEEE Trans. Signal Process.* **43**, 188 (1995).
- [47] A. J. Barbour and R. L. Parker, psd: Adaptive, sine multitaper power spectral density estimation for R, *Comput. Geosci.* **63**, 1 (2014).
- [48] F. Harris, On the use of windows for harmonic analysis with the discrete Fourier transform, *Proc. IEEE* **66**, 51 (1978).
- [49] bSPOD—Band-ensemble spectral proper orthogonal decomposition, <https://github.com/jakobvonsaldern/bspod>.
- [50] L. Heidt and T. Colonius, Optimal frequency resolution for spectral proper orthogonal decomposition, [arXiv:2402.15775](https://arxiv.org/abs/2402.15775).
- [51] Y. Zhang, L. N. Cattafesta, and L. Ukeiley, Spectral analysis modal methods (SAMMs) using non-time-resolved PIV, *Exp. Fluids* **61**, 226 (2020).
- [52] A. V. G. Cavalieri, D. Rodríguez, P. Jordan, T. Colonius, and Y. Gervais, Wavepackets in the velocity field of turbulent jets, *J. Fluid Mech.* **730**, 559 (2013).