

Towards real-time reconstruction of velocity fluctuations in turbulent channel flow

Rahul Arun ,* H. Jane Bae , and Beverley J. McKeon [†]*Graduate Aerospace Laboratories, California Institute of Technology, Pasadena, California 91125, USA*

(Received 18 January 2023; accepted 1 June 2023; published 27 June 2023)

We develop a framework for efficient streaming reconstructions of turbulent velocity fluctuations from limited sensor measurements with the goal of enabling real-time applications. The reconstruction process is simplified by computing linear estimators using flow statistics from an initial training period and evaluating their performance during a subsequent testing period with data obtained from direct numerical simulation. We address cases where (i) no, (ii) limited, and (iii) full-field training data are available using estimators based on (i) resolvent modes, (ii) resolvent-based estimation, and (iii) spectral proper orthogonal decomposition modes. During training, we introduce blockwise inversion to accurately and efficiently compute the resolvent operator in an interpretable manner. During testing, we enable efficient streaming reconstructions by using a temporal sliding discrete Fourier transform to recursively update Fourier coefficients using incoming measurements. We use this framework to reconstruct with minimal time delay the turbulent velocity fluctuations in a minimal channel at $Re_\tau \approx 186$ from sparse planar measurements. We evaluate reconstruction accuracy in the context of the extent of data required and thereby identify potential use cases for each estimator. The reconstructions capture large portions of the dynamics from relatively few measurement planes when the linear estimators are computed with sufficient fidelity. We also evaluate the efficiency of our reconstructions and show that the present framework has the potential to help enable real-time reconstructions of turbulent velocity fluctuations in an analogous experimental setting.

DOI: [10.1103/PhysRevFluids.8.064612](https://doi.org/10.1103/PhysRevFluids.8.064612)

I. INTRODUCTION

A. Real-time estimation and control

Processing measurements generated by turbulent flows in real time is an increasingly important problem in applications including flow estimation and control. As depicted in Fig. 1, estimates can be used to inform control schemes in real-time applications. Flow estimation problems may be further subdivided into smoothing, filtering, and prediction problems. These problems aim to inform estimates of past, present, and future states (respectively) using incoming measurements that are usually incomplete and noisy. Real-time predictions can be useful for active control, but they are typically more challenging than (smoothing or filtering) reconstructions of turbulent flows due to multiscale dynamics and sensitivity to initial conditions.

Real-time flow estimation tasks are often practically limited by low-fidelity measurements and task-related time constraints, e.g., in the context of aviation meteorology [1]. Measurements can be improved by optimizing the design [2] and configuration [3] of sensors. Computational processing

*rarun@caltech.edu

[†]Present address: Department of Mechanical Engineering, Stanford University, Stanford, California 94305, USA.

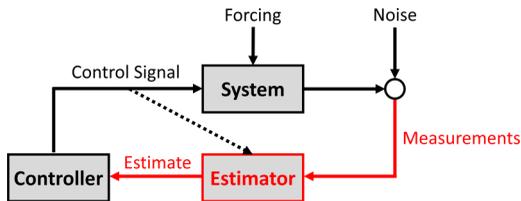


FIG. 1. High-level block diagram depicting a framework in which a real-time control scheme is informed by real-time estimates generated by noisy measurements. We focus on the real-time estimation problem (red) in the context of flow reconstruction.

can be expedited by leveraging data-driven and physics-based methods [4] that often reduce the dimensionality of the problem or introduce a simplifying model. In the context of turbulent jets, Sasaki *et al.* [5] demonstrated that the parabolized stability equations can be used to estimate (in real time) the pressure at downstream sensors using upstream pressure measurements.

Beyond estimation, real-time control schemes require yet further empirically or heuristically prescribed simplifications to ensure in-time actuation. Maia *et al.* [6] used a reactive control scheme to experimentally attenuate centerline velocity fluctuations in a forced jet ($Re_D = 50\,000$) by generating wave packets that destructively interfere with stochastic disturbances generated at the nozzle. Abbassi *et al.* [7] demonstrated that actuating an array of cross-flowing jets that penetrate into the log region of a turbulent boundary layer ($Re_\tau = 14\,400$) can influence the resulting skin friction drag and turbulence intensity.

Previous real-time flow estimation and control schemes demonstrate that, by combining governing equations with measurements and simplifying models, hybrid methods allow for a tailored balance between the efficacy and the speed of the scheme. One commonality of these methods is that their simplifying assumptions are often not well justified from first principles. In the present investigation, we focus on the problem of real-time flow reconstruction in the context of wall-bounded turbulence. Even for this relatively simple estimation problem, the task of efficiently and accurately representing turbulent dynamics from first principles [i.e., the Navier-Stokes equations (NSE)] using minimal assumptions has remained challenging.

B. Reconstruction techniques

Reduced-order models (ROMs) often provide efficient flow representations to reduce the complexity of high-dimensional turbulent flows, e.g., using proper orthogonal decomposition (POD) [8] and dynamic mode decomposition (DMD) [9]. For statistically stationary flows, spectral POD (SPOD) provides an efficient means of identifying coherent structures that retains a direct relationship to resolvent analysis and DMD [10]. Unless otherwise stated, the SPOD modes we refer to are those of the velocity fluctuations. Truncated SPOD mode estimation (TSME) often provides an efficient means of flow reconstruction since SPOD modes form an optimal orthogonal basis in terms of the variance captured by a given subset of modes [11]. Data-driven methods like SPOD are often limited in that they neglect the governing equations and require extensive postprocessing. However, Ghate *et al.* [12] showed that supplementing SPOD-based truncations with physics-based enrichment using Gabor modes enables representations of flows over a broad range of scales. Moreover, SPOD is amenable to a streaming formulation [13] and convolution-based strategies for time-domain analysis [14], thereby reducing its computational requirements.

Equation-based frameworks have been developed in conjunction with data-driven methods to efficiently estimate turbulent flows. Techniques from control theory are often used for flow estimation in the time domain (e.g., Kalman filters [15–17]) and the frequency domain (e.g., Wiener filters [18]). Techniques like linear stochastic estimation (LSE) [19] and its spectral variant [20] are rooted in conditional estimation. LSE-based techniques are often used with data-driven [21,22]

and equation-based [23,24] models to augment flow reconstructions. Other methods use simplified governing models to produce forward and backward estimates that augment reconstructions from low-resolution and multiresolution temporal measurements [25,26]. One central challenge in these estimation techniques is addressing the nonlinearity of the governing equations.

Resolvent analysis [27,28] provides a powerful framework for addressing the nonlinearity of the NSE and identifying energetic linear amplification mechanisms using minimal assumptions. In a manner related to the work of Farrell and Ioannou [29], the turbulent fluctuation dynamics are recast as a linear system forced by their nonlinearity in the resolvent framework. The resolvent operator often admits a low-rank representation, and it can be constructed using a base flow profile, which can be modeled or learned from data, e.g., via data assimilation [30,31].

One flavor of flow reconstruction using resolvent analysis is truncated response mode estimation (TRME). TRME revolves around the singular value decomposition (SVD) of the resolvent operator, which produces orthonormal bases for the forcing and response that are related through corresponding gains or singular values. Modal truncation is performed based on the (often) low-rank nature of the resolvent operator, which occurs when the gains associated with a small number of leading modes dominate those of the remaining (suboptimal) modes. Moarref *et al.* [32] demonstrated that streamwise energy amplification in each half of a turbulent channel ($Re_\tau = 2003$) is well captured by a rank-1 approximation of the resolvent operator. Moarref *et al.* [33] then used convex optimization techniques to compute resolvent mode weights, which encode dynamical information, using a similar low-rank approximation to capture the velocity spectra in a turbulent channel. In similar fashions, TRME was used to approximate energetic structures in flow around a cylinder [31] and in lid-driven cavity flow ($Re = 1200$) with a known (2D) mean flow [34]. In all of these cases, the resolvent modes are weighted using spatially isolated velocity measurements in either Fourier or physical space. In such setups, it is important that the measurements at least partially capture the most energetic regions of the flow. For example, Symon [30] showed that tailoring probe locations to capture energetic regions of both low-frequency wake modes and high-frequency shear layer modes significantly improves reconstruction accuracy. Beneddine *et al.* [35] found a similar result by separately capturing low-frequency and high-frequency regions in a backward-facing step flow configuration. Recently, a quasisteady resolvent analysis has been used to reconstruct high-frequency fluctuations using the phase of a lower-frequency periodic motion [36]. Resolvent truncations have also recently been used to design H_2 -optimal estimators and controllers [37] and select optimal sensor and actuator locations [38] in cylinder flows.

While TRME provides an efficient and informative approximation, the assumption of a fixed-rank [39] truncation may not be appropriate across all frequencies. In contrast to TRME, resolvent-based estimation (RBE) revolves around the SVD of an input-output operator that modifies the resolvent operator to reflect sensor locations. Since flow sensors are usually limited and sparsely spaced, RBE has shown promise in producing reconstructions of both flow statistics [40] and dynamics [41] from limited data without truncating the resolvent operator *a priori*. RBE [40] relates response measurement statistics to observable forcing statistics via an input-output operator and (typically) models the unobservable forcing as uncorrelated with the observable forcing. The estimated forcing statistics then yield estimates of the response statistics via the resolvent operator. Building from this approach, Martini *et al.* [41] derived optimal, noncausal transfer functions relating the measurements to the full response when the second-order forcing statistics and resolvent operator are known. Further, Morra *et al.* [42] showed that the statistics in a turbulent channel can be reasonably modeled using a low-rank SPOD approximation of the forcing. These studies reveal that the tradeoff between accurately capturing the forcing color and retaining a sparse measurement configuration can be somewhat alleviated by modeling the structure of the nonlinearity. For example, in turbulent channel flows, including an eddy-viscosity model [17] can improve the accuracy of RBE techniques for the response [43] and its statistics [40,42,44] with minimal input data.

We primarily focus on noncausal forms of linear estimators, including that derived for optimal resolvent-based estimation (ORBE) [41,43]. More recently, a causal resolvent-based formulation for estimation and control [45] was developed via the Wiener-Hopf formalism. However, the

numerical machinery for forming this causal estimator is typically more complex than that of the noncausal estimator due to an additional term in the estimation problem. By instead applying the noncausal formulation to the real-time reconstruction problem, we avoid these complex methods and retain a relatively simple framework. Noncausal estimation techniques (including ORBE) are not typically amenable to real-time applications [43], but we circumvent this limitation by continuously updating the estimated temporal Fourier coefficients over a small window at every time step. Further, using this sliding window enables efficient recursive techniques for updating the coefficients [46]. The implementation of this technique, which has linear algorithmic complexity, differs from the streaming Fourier sums technique of Schmidt and Towne [13], which has quadratic algorithmic complexity but a smaller memory footprint.

Nonlinear estimation techniques, especially those incorporating machine learning techniques, can also reconstruct flows with competitive accuracy. For example, Fukami *et al.* [47] reported relative errors of less than 15% for spatial superresolution of a channel flow using a coarse input with six elements across each half of the channel. Although they achieve significant data compression, the coarsened input for estimation, obtained via averaging, is not conducive to typical experimental measurements. By contrast, our investigation focuses on practical considerations for experimental implementation, including time delays, measurement configurations, scalability, and uncertainty quantification. Working in the linear estimation setting is further advantageous in that it is amenable to a wealth of estimation and control theory for linear systems. Further, in contrast to many machine learning techniques, the present linear estimation techniques are interpretable in terms of their optimality and through the basis and flow statistics imposed by the linear model. Finally, whereas machine learning techniques can take significantly longer (e.g., multiple days [47]) to train estimators on optimized hardware, we introduce estimators that can be constructed from training data within an hour on a laptop.

C. Contributions

The goal of this study is to construct a method for reconstruction of turbulent velocity fluctuations from sparse measurements, adding to the vast flow estimation literature. For this, we develop a new framework to enable efficient streaming capabilities with the potential for real-time implementation. In Sec. II we detail the governing equations and formulate the statistical framework, based on the generalized Wiener filter, that we use to reconstruct the fluctuations. We also detail the application of methods inspired by TRME, ORBE, and TSME to this more general framework. In Sec. III, we introduce efficient methods that enable streaming flow reconstructions with minimal time delay and improved scalability compared to standard techniques. When forming linear estimators from training data, we introduce blockwise inversion to efficiently compute the resolvent operator. When reconstructing the flow from testing data, we apply a recursive sliding discrete Fourier transform (SDFT) to efficiently update the inputs to the estimators in a streaming fashion. In Sec. IV we evaluate the performance of the reconstructions using each method. We specifically address (i) the fidelity of the training data and the testing data, (ii) the validity of the models underlying the estimators, (iii) the potential use cases for each reconstruction method, and (iv) the reconstruction efficiency in the context of real-time applications. In doing so, we demonstrate that practical reconstruction speeds can be achieved using modest computational resources while retaining competitive accuracies with other linear estimation techniques. Finally, we summarize our reconstruction methods, the performance of the estimators, and promising future prospects and applications in Sec. V.

II. BACKGROUND METHODOLOGY

A. Governing equations

We focus on turbulent channel flow between two parallel walls, for which we denote the streamwise, wall-normal, and spanwise directions by x , y , and z and the corresponding velocities by \mathcal{U} , \mathcal{V} ,

and \mathcal{W} , respectively. This flow is governed by the nondimensional, incompressible Navier-Stokes equations (NSE)

$$\partial_t \mathbf{U} + (\mathbf{U} \cdot \nabla) \mathbf{U} + \nabla \mathcal{P} = \frac{1}{\text{Re}_\tau} \nabla^2 \mathbf{U}, \quad \nabla \cdot \mathbf{U} = 0, \quad (1)$$

where $\mathbf{U} = [U, V, W]^T$ is the velocity vector, \mathcal{P} is pressure, and t denotes time. For channel flow, $\text{Re}_\tau = u_\tau h / \nu$ is defined in terms of the channel half-height h , the friction velocity u_τ , and the kinematic viscosity ν . Lengths, velocities, and pressures are normalized by h , u_τ , and ρu_τ^2 , respectively, where ρ is the fluid density.

We combine the velocity and pressure into a single state variable, $\mathbf{Q} = [U, V, W, \mathcal{P}]^T$, whose evolution is governed by the NSE. The state fluctuations are expressed by (Reynolds) decomposing the state vector as $\mathbf{Q} = \mathbf{Q} + \mathbf{q}$, where \mathbf{Q} and \mathbf{q} represent the base flow and the fluctuations, respectively. Here we specify the base flow as the mean state profile, $\mathbf{Q}(y) = \overline{\mathbf{Q}}(y) = [U(y), V(y), W(y), P(y)]^T$, where $V(y) = W(y) = 0$ for channel flow [48]. The corresponding fluctuations are Fourier transformed in time and in the spatially homogeneous and periodic directions (x and z here), giving

$$\hat{\mathbf{q}}(y, \mathbf{k}) = \begin{bmatrix} \hat{\mathbf{u}}(y, \mathbf{k}) \\ \hat{p}(y, \mathbf{k}) \end{bmatrix} = \mathcal{F}\{\mathbf{q}(\mathbf{x}, t)\} = \frac{1}{8\pi^3} \iiint_{-\infty}^{\infty} \mathbf{q}(\mathbf{x}, t) e^{i(\omega t - k_x x - k_z z)} dt dx dz, \quad (2)$$

where $i = \sqrt{-1}$, $\mathbf{x} = [x, y, z]^T$ and each triplet, $\mathbf{k} = [k_x, k_z, \omega]^T$, contains the streamwise wave numbers, spanwise wave numbers, and (angular) temporal frequencies, respectively. The transformed differential operators are given by $\hat{\nabla} = [ik_x, \partial_y, ik_z]^T$ and $\hat{\nabla}^2 = \partial_{yy} - k^2$, where $k^2 = k_x^2 + k_z^2$.

Applying the Reynolds decomposition and the Fourier transform defined in (2), (1) can be used to express the state fluctuation dynamics for each triplet via a linear subsystem that is externally forced by the nonlinear advection terms. Correspondingly, each linear subsystem may be written as

$$\mathcal{L}(\mathbf{k}) \hat{\mathbf{q}}(y, \mathbf{k}) = \begin{bmatrix} \mathcal{L}_B(\mathbf{k}) & \hat{\nabla} \\ \hat{\nabla}^T & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{u}}(y, \mathbf{k}) \\ \hat{p}(y, \mathbf{k}) \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{f}}(y, \mathbf{k}) \\ 0 \end{bmatrix}, \quad (3)$$

where $\mathbf{k} \neq [0, 0, 0]^T$. Here \mathcal{L}_B represents the ‘‘basic’’ terms in the unsteady momentum equations, which are linear in the *velocity* fluctuations, and $\hat{\mathbf{f}}(y, \mathbf{k}) = \mathcal{F}\{-(\mathbf{u} \cdot \nabla) \mathbf{u} + \overline{(\mathbf{u} \cdot \nabla) \mathbf{u}}\}$ represents the nonlinear forcing. While the subsystem for each triplet is *externally* forced by velocity fluctuations satisfying $\mathbf{k}' \neq \mathbf{k}$, the system is *internally* forced by its nonlinearity. In Sec. II B we discretize this state-space formulation of the Navier-Stokes system and express the corresponding resolvent operator in discrete form.

B. Discretized linear system

When formed directly from the continuous governing equations, the linear subsystems in (3) are infinite-dimensional. However, by discretizing the equations onto a grid, they assume a finite-dimensional representation based on the spatial resolution in the wall-normal direction. As depicted in Fig. 2, the direct numerical simulation (DNS) data in the present case are staggered [49] on a Cartesian grid. For a wall-normal domain discretized into N_y cells, the discrete representations of $\hat{\mathbf{u}}$ and $\hat{\mathbf{q}}$ are column vectors of sizes $N_u = 3N_y + 1$ and $N_q = N_u + N_y$, respectively. The numerical quadrature (i.e., wall-normal integration) weights associated with center and edge quantities are Δy_j and $0.5(\Delta y_{j-1} + \Delta y_j)$, respectively.

Since state variables are Fourier transformed in x and z , we accurately reference the cell-center values for $\hat{\mathbf{u}}$ and $\hat{\mathbf{w}}$ by shifting their Fourier transform origins by half a cell width in x and z , respectively. To retain a direct analogy with the (discrete) governing equations used in simulating the flow, we maintain \hat{v} and the mean shear, $\partial_y U$, at the wall-normal edges, y_j^e and y_{j+1}^e , of each cell. Correspondingly, we represent the mean shear using a bidiagonal matrix, $\hat{\partial}_y U$, that acts to average

center. As such, it may be expressed as

$$S(\mathbf{k}) = \begin{bmatrix} S_x(\mathbf{k})\mathbf{I}_c & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_e & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & S_z(\mathbf{k})\mathbf{I}_c \end{bmatrix}, \quad S_x = \exp\left(\frac{ik_x^* \Delta x}{2}\right), \quad S_z = \exp\left(\frac{ik_z^* \Delta z}{2}\right), \quad (8)$$

where $S^{-1} = S^H$ and $(\cdot)^H$ is the conjugate transpose. The input matrix,

$$B = \begin{bmatrix} \mathbf{I}_c & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_e & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_c \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (9)$$

is used to isolate the forcing to the momentum equations. For convenience, we also use it to extract the velocity fluctuations from the state vector. For each triplet, the velocity fluctuations may then be expressed as a linear response to their nonlinear forcing via

$$\hat{\mathbf{u}}(\mathbf{k}) = \mathbf{B}^T \hat{\mathbf{q}}(\mathbf{k}) = \mathbf{R}_u(\mathbf{k}) \hat{\mathbf{f}}(\mathbf{k}), \quad (10)$$

where

$$\mathbf{R}_u = (\mathbf{B}\mathbf{S})^H \mathbf{L}^{-1} (\mathbf{B}\mathbf{S}) \quad (11)$$

is the one-dimensional resolvent [50] operator. We directly enforce the no-penetration boundary condition in computing $\mathbf{L}^{-1} \mathbf{B}\mathbf{S}$.

C. Flow reconstruction

We formulate the flow reconstruction problems for TRME (Sec. IIC 1), ORBE (Sec. IIC 2), and TSME (Sec. IIC 3) in triplet space. Correspondingly, our goal is to reconstruct velocity fluctuations in the turbulent channel using spatially sparse measurements in the wall-normal direction. We therefore consider a collection of N_{plane} spatially isolated planar observations,

$$\hat{\mathbf{y}}_n(\mathbf{k}) = \mathbf{C}(\mathbf{k}) \hat{\mathbf{u}}(\mathbf{k}) + \hat{\mathbf{n}}(\mathbf{k}) = \hat{\mathbf{y}}(\mathbf{k}) + \hat{\mathbf{n}}(\mathbf{k}), \quad (12)$$

where $\mathbf{C} \in \mathbb{R}^{3N_{\text{plane}} \times N_u}$ is the observation matrix, $\hat{\mathbf{n}}$ represents the measurement noise, and $\hat{\mathbf{y}} = \mathbf{C}\hat{\mathbf{u}}$ isolates the noiseless component of the observations. For all triplets of interest, \mathbf{C} isolates *collocated* velocity fluctuation measurements at cell-centered wall-parallel planes by averaging adjacent wall-normal velocities. The noise is assumed to represent a zero-mean random process that is uncorrelated with the flow. In this formulation, $\hat{\mathbf{n}}$ captures errors resulting from employing a discrete Fourier transform over a finite temporal window and any other sources of contamination. In an experimental setting, observations of the form in (12) may be generated by (multiplane) wall-parallel particle image velocimetry (PIV) measurements.

We formulate optimal linear transfer functions by operating under the generalized Wiener filter formalism. For each reconstruction method, we consider generic linear models for the flow and observations of the form

$$\hat{\mathbf{u}} = \mathbf{H}_u \mathbf{b}, \quad \hat{\mathbf{y}}_n = \mathbf{H}_y \mathbf{b} + \hat{\mathbf{n}}, \quad (13)$$

where $\mathbf{H}_y = \mathbf{C}\mathbf{H}_u$ and \mathbf{b} is a set of unknown coefficients. In weighted form, (13) may be written as

$$\mathbf{W}_u \hat{\mathbf{u}} = \mathbf{H}_u^w \mathbf{W}_b \mathbf{b}, \quad \mathbf{W}_y \hat{\mathbf{y}}_n = \mathbf{H}_y^w \mathbf{W}_b \mathbf{b} + \mathbf{W}_y \hat{\mathbf{n}}, \quad (14)$$

where $\mathbf{H}_u^w = \mathbf{W}_u \mathbf{H}_u \mathbf{W}_b^{-1}$, $\mathbf{H}_y^w = \mathbf{C}_w \mathbf{H}_u^w$, and $\mathbf{C}_w = \mathbf{W}_y \mathbf{C} \mathbf{W}_u^{-1}$. Here we have assumed \mathbf{W}_y , \mathbf{W}_u , and \mathbf{W}_b are Hermitian, positive-definite weighting matrices such that the norms of the unweighted vectors corresponding to their respective weighted inner products are equivalent to the Euclidean 2-norms of the weighted vectors. Unless otherwise stated, all weightings in the present investigation are real,

diagonal matrices that contain the square roots of the numerical quadrature weights to ensure that the relevant norms reproduce energies integrated over the wall-normal direction.

Given a linear model, \mathbf{H}_u , and an observation matrix, \mathbf{C} , the generalized Wiener filter provides an optimal linear estimate of the coefficients, $\tilde{\mathbf{b}}$, assuming that the second-order statistics of the coefficients, $\mathbf{S}_{bb} = \mathbb{E}(\mathbf{b}\mathbf{b}^H)$, and the noise, $\mathbf{S}_{nn} = \mathbb{E}(\hat{\mathbf{n}}\hat{\mathbf{n}}^H)$, or models thereof are known. Optimality is typically defined in terms of minimizing the mean-squared error of the estimated coefficients. However, Martini *et al.* [41] showed (in the context of RBE) that the corresponding optimal estimator is the stationary point of the entire coefficient error CSD, not just its trace. Furthermore, this stationary point also applies to the velocity fluctuation error CSD, highlighting the correspondence between optimal estimates of the coefficients and the flow. Hence, for each reconstruction method, we express the estimator via a linear transfer function, \mathbf{T}_u , that relates the noisy measurements to the flow estimates as

$$\tilde{\mathbf{b}} = \mathbf{T}_b \hat{\mathbf{y}}_n, \quad \tilde{\mathbf{u}} = \mathbf{T}_u \hat{\mathbf{y}}_n, \quad (15)$$

where $\mathbf{T}_u = \mathbf{H}_u \mathbf{T}_b$. Hereafter, unless otherwise stated, all quantities denoted as $(\tilde{\cdot})$ represent estimates/models of true flow quantities or the building blocks thereof.

The transfer function may be explicitly written as

$$\mathbf{T}_u = \mathbf{S}_{uy} \mathbf{S}_{yy,n}^{-1} = \mathbf{H}_u \mathbf{S}_{bb} \mathbf{H}_y^H (\mathbf{H}_y \mathbf{S}_{bb} \mathbf{H}_y^H + \mathbf{S}_{nn})^{-1}, \quad (16)$$

where $\mathbf{S}_{yy,n} = \mathbb{E}(\hat{\mathbf{y}}_n \hat{\mathbf{y}}_n^H)$ and we have used that the noise is uncorrelated with the flow such that $\mathbf{S}_{un} = \mathbf{0}$. Absent knowledge of the noise CSD, we model it as $\tilde{\mathbf{S}}_{nn} = \epsilon \mathbf{I}$ (for small ϵ) to ensure that $\mathbf{S}_{yy,n}^{-1}$ remains numerically well posed while approaching the zero-noise limit. The transfer function, \mathbf{T}_u , applies to both (13) and (14), demonstrating that it is independent of the weightings placed on the response, coefficients, and measurements/noise. We hereafter refer to the linear models via (13), but note that the weightings may still be incorporated implicitly in the models of \mathbf{H}_u , \mathbf{S}_{bb} , and/or \mathbf{S}_{nn} .

Importantly, (13) is a model in the sense that $\mathbf{H}_u \mathbf{b}$ may be an ill-posed representation of the flow. Therefore, in addition to the knowledge or models of \mathbf{S}_{bb} and \mathbf{S}_{nn} , the well-posedness of the assumed linear model, \mathbf{H}_u , is also implicit to the optimality of the linear estimator produced by the generalized Wiener filter. All methods we consider incorporate the estimated mean flow in learning \mathbf{H}_u from training data, and TSME and ORBE further incorporate the response and (thereby) the nonlinear forcing CSDs, respectively. In what follows, we derive expressions for the transfer functions corresponding to each reconstruction method. We discuss our methodology of efficiently applying the transfer functions to streaming measurements in Sec. III.

1. Truncated response mode estimation

TRME is an equation-based reconstruction framework that requires prior knowledge of the mean profile, but not of the second-order statistics. Therefore, TRME requires minimal data collection and modeling to form an estimator, but it is founded on the assumptions that a low-rank approximation of the resolvent operator can capture the relevant dynamics and that the corresponding forcing (statistics) may be modeled from first principles.

The low-rank approximation of the weighted resolvent operator is expressed by truncating its SVD as

$$\mathbf{R}_u^w = (\mathbf{W}_u \mathbf{\Psi}) \mathbf{\Sigma} (\mathbf{W}_f \mathbf{\Phi})^H \approx (\mathbf{W}_u \tilde{\mathbf{\Psi}}) \tilde{\mathbf{\Sigma}} (\mathbf{W}_f \tilde{\mathbf{\Phi}})^H = \tilde{\mathbf{R}}_u^w, \quad (17)$$

where

$$\tilde{\mathbf{\Psi}} = \mathbf{\Psi} \tilde{\mathbf{M}}, \quad \tilde{\mathbf{\Phi}} = \mathbf{\Phi} \tilde{\mathbf{M}}, \quad \tilde{\mathbf{\Sigma}} = \tilde{\mathbf{M}}^T \mathbf{\Sigma} \tilde{\mathbf{M}}, \quad \tilde{\mathbf{M}} = [\tilde{\mathbf{I}} \quad \mathbf{0}]^T, \quad (18)$$

and $\tilde{\mathbf{I}}$ is the identity matrix of size $N_{\text{mode}} \leq N_u$. Here the forcing ($\phi_j \in \mathbb{C}^{N_u}$) and response ($\psi_j \in \mathbb{C}^{N_u}$) modes are the columns of $\mathbf{\Phi}$ and $\mathbf{\Psi}$, respectively, and their weightings ensure that orthonormality

enforces an energy norm. These pairs of modes are ranked according to their gains, σ_j , which are the diagonal elements of Σ and ordered such that $\sigma_j \geq \sigma_{j+1}$.

The approximation in (17) is founded on the premise that the gains for modes $j \leq N_{\text{mode}}$ dominate those for modes $j > N_{\text{mode}}$. However, response modes associated with relatively small gains can remain important to the velocity fluctuations if the forcing has a large projection onto the corresponding forcing modes. For each triplet, we collect the gain-scaled projections of the nonlinear forcing onto the N_{mode} selected forcing modes into a column vector of mode weights,

$$\mathbf{c} = \tilde{\Sigma}(\mathbf{W}_f \tilde{\Phi})^H (\mathbf{W}_f \hat{\mathbf{f}}), \quad (19)$$

that characterizes the contributions of the selected response modes to the response. The linear model corresponding to the truncated flow representation is given by $\hat{\mathbf{u}} \approx \tilde{\Psi} \mathbf{c}$.

The SVD of the resolvent operator does not provide information about \mathcal{S}_{cc} without further assumptions. However, we may form a model of this CSD by recalling that resolvent modes and SPOD modes are identical when the resolvent mode weights are uncorrelated [10]. To achieve this, we assume an uncorrelated weighted forcing, $\mathbf{W}_f \mathbb{E}(\hat{\mathbf{f}} \hat{\mathbf{f}}^H) \mathbf{W}_f^H = \mathbf{I}$, for the present TRME-inspired estimator. Noting the orthonormality of the forcing modes, the modeled mode weight CSD is simply $\tilde{\mathcal{S}}_{cc} = \tilde{\Sigma} \tilde{\Sigma}^H = \tilde{\Sigma}^2$. Since this statistical perspective implies that resolvent modes form an approximation of SPOD modes, we expect that a truncated SPOD mode basis is more well posed to fitting the response than a truncated (resolvent) response mode basis.

Assuming the resolvent response modes form a suitable basis for the linear model, we may form the constituents of the estimation transfer function in (16) by taking

$$\mathbf{H}_u = \tilde{\Psi}, \quad \mathbf{H}_y = \tilde{\Psi}_y, \quad \text{and} \quad \mathbf{b} = \mathbf{c}, \quad (20)$$

where $\tilde{\Psi}_y = \mathbf{C} \tilde{\Psi}$. With these definitions, the transfer function is expressed as

$$\mathbf{T}_u^{\text{TRME}} = \tilde{\Psi} \tilde{\Sigma}^2 \tilde{\Psi}_y^H (\tilde{\Psi}_y \tilde{\Sigma}^2 \tilde{\Psi}_y^H + \tilde{\mathcal{S}}_{nn})^{-1}. \quad (21)$$

This estimation method contrasts slightly from the direct linear least-squares method of Gómez *et al.* [34] in that we assume that the coefficient statistics are known (and uncorrelated). However, since this method arises from the generalized Wiener filter formulation, employing a suitable model on $\tilde{\mathcal{S}}_{cc}$ can mitigate the instability of higher-order modal truncations associated with conventional TRME [41]. Hence, this method provides an alternative to limiting modal truncations such that $N_{\text{mode}} < N_{\text{plane}}$, which discards sensor information.

2. Optimal resolvent-based estimation

The TRME-based method in Sec. IIC 1 uses an (unrealistic) uncorrelated forcing model to produce estimators that only require a mean profile. ORBE [41] provides a practical means of improving the forcing model by estimating its CSD using limited auxiliary data. In ORBE, the constituents of the linear estimator are given by

$$\mathbf{H}_u = \mathbf{R}_u, \quad \mathbf{H}_y = \mathbf{R}_y, \quad \text{and} \quad \mathbf{b} = \hat{\mathbf{f}}, \quad (22)$$

where we refer to the input-output operator, $\mathbf{R}_y = \mathbf{C} \mathbf{R}_u$, distinctly from the resolvent operator. Hence, the ORBE transfer function is given by

$$\mathbf{T}_u^{\text{ORBE}} = \mathbf{R}_u \tilde{\mathcal{S}}_{ff} \mathbf{R}_y^H (\mathbf{R}_y \tilde{\mathcal{S}}_{ff} \mathbf{R}_y^H + \tilde{\mathcal{S}}_{nn})^{-1}. \quad (23)$$

Even though ORBE was originally formulated using an unweighted inner product [41], (23) and (14) reinforce that the Wiener filter origins of the transfer function imply it is invariant to nondegenerate weights imposed on the input-output formulation. We emphasize that (23) is an optimal noncausal transfer function and that forming the optimal causal transfer functions derived via the Wiener-Hopf formalism [45] requires more complex numerical methods.

Instead of truncating a modal expansion of the dynamics *a priori*, ORBE identifies important parts of the forcing even if they reside in linearly suboptimal resolvent modes, e.g., when their projections onto the relevant forcing bases are large. Further, while the truncations associated with TRME and TSME may result in an ill-posed basis for representing the full fluctuation dynamics, ORBE retains a full-rank (and therefore well-posed) representation of the velocity fluctuations. When all resolvent modes are retained, the TRME-based method discussed in Sec. II C 1 is equivalent to ORBE under the assumption of an uncorrelated weighted forcing. It is well known that this assumption is not typically valid for turbulent flows, including channel flow, since the forcing color contributes significantly to the forcing structure, e.g., via destructive interference [42,51]. Correspondingly, the focus of ORBE is identifying a suitable ROM of the forcing statistics that enables accurate reconstructions.

Following previous studies [40,41,45], we estimate the forcing statistics using an auxiliary set of measurements, \hat{y}'_n , with corresponding CSD, $S'_{yy,n}$, observation matrix, C' , and input-output operator, $R_{y'} = C'R_u$. The forcing estimate is given by

$$\tilde{S}_{ff} = R_{y'}^{\dagger,n} S'_{yy} (R_{y'}^{\dagger,n})^H, \quad R_{y'}^{\dagger,n} = R_{y'}^H (R_{y'} R_{y'}^H + \tilde{S}'_{nn})^{-1}, \quad (24)$$

where $R_{y'}^{\dagger,n}$ maps the auxiliary observations to the smallest regularized least-squares forcing producing them and reduces to a pseudoinverse in the limit of vanishing noise. We approach this limit by using the same noise model, $\tilde{S}'_{nn} = \epsilon I$, as that used for the reconstruction observations. One advantage of this estimation technique is that a limited set of measurements can produce a global forcing estimate, whereas computing SPOD modes (see Sec. II C 3) typically requires knowledge of the full-field response statistics.

When full-field statistics are available, the ORBE transfer function reduces to the generalized Wiener filter transfer function, $S_{uy} S_{yy,n}^{-1}$, which is independent of the governing equations. With noisy or incomplete statistics, ORBE implicitly assumes that the coefficients are the resolvent forcing. This distinction is important since the forcing statistics may not be an optimal setting for approximating $S_{uy} S_{yy,n}^{-1}$ from limited data. Hence, while ORBE provides a forcing-based means of flow reconstruction, determining the optimal coefficient coordinates for flow reconstruction remains an open question.

3. Truncated SPOD mode estimation

TSME provides a framework for estimating velocity fluctuations by truncating the SPOD basis to efficiently represent coherent structures in the flow. TSME requires the mean profile and second-order statistics of the velocity (response) fluctuations, but it does not consider the dynamics associated with the governing equations. For each triplet, we form the response CSD [52], S_{uu} , using an ensemble of N_{real} realizations of the flow. The corresponding SPOD modes, $\theta_j \in \mathbb{C}^{N_u}$, are given by the columns of Θ and computed via a direct eigenvalue problem,

$$S_{uu} (W_u^H W_u) \Theta = \Theta \Lambda. \quad (25)$$

As before, the weighting matrix, W_u , ensures that orthonormality enforces an energy norm. This formulation is typically employed when $N_{\text{real}} \geq N_u$, whereas the method of snapshots is standard when $N_{\text{real}} < N_u$. Since it is often the case that $N_{\text{real}} \ll N_u$ when estimating turbulent flows, the method of snapshots is often less computationally expensive than the direct eigenvalue problem. However, since the CSD sizes are not prohibitive in the present investigation, we simply employ the direct eigenvalue formulation in (25).

The matrix Λ contains the eigenvalues that rank the SPOD modes in terms of the variance they capture in the response. Since a given subset of SPOD modes captures more energy than in any other orthogonal basis [11], truncating the SPOD mode basis provides a natural means of efficiently representing dominant coherent flow structures. The truncated representation of the flow is given by

TABLE I. Relevant parameters for data sampled from the DNS, where nondimensional parameters with superscripts $(\cdot)^+$ are normalized in viscous units (e.g., $\Delta t^+ = \text{Re}_\tau \Delta t$). The values of u_τ , ν , and h are dimensional, but the dimensional units are arbitrary so long as they are consistent.

Δx^+	Δz^+	Δy_{\min}^+	Δy_{\max}^+	Δt^+	u_τ	ν	h
10.3	5.16	0.172	7.58	0.532	0.0572	0.000308	1.00

$\hat{\mathbf{u}} \approx \tilde{\Theta} \mathbf{a}$, where

$$\tilde{\Theta} = \Theta \tilde{\mathbf{M}}, \quad \tilde{\Lambda} = \tilde{\mathbf{M}}^T \Lambda \tilde{\mathbf{M}}, \quad (26)$$

and $\tilde{\mathbf{M}}$ again selects the first N_{mode} modes of the expansion (for $N_{\text{mode}} \leq N_{\text{real}}$).

The SPOD mode expansion is advantageous in that the coefficients are uncorrelated and the relevant CSD is given by $\mathbf{S}_{aa} = \mathbf{a} \mathbf{a}^H = \tilde{\Lambda}$. As such, the elements of the transfer function in (16) are given by

$$\mathbf{H}_u = \tilde{\Theta}, \quad \mathbf{H}_y = \tilde{\Theta}_y, \quad \text{and} \quad \mathbf{b} = \mathbf{a}, \quad (27)$$

where $\tilde{\Theta}_y = \mathbf{C} \tilde{\Theta}$. Correspondingly, the SPOD modes and energies suffice in providing both the linear model and the coefficient statistics implicit to the optimality of the linear transfer function. When all modes are retained, the TSME estimator converges to the generalized Wiener filter, $\mathbf{S}_{uy} \mathbf{S}_{yy,n}^{-1}$. Even when no modal truncation is performed, it is important to capture sufficiently many realizations to ensure the SPOD modes provide a well-posed basis for the flow.

For TSME, the transfer function is given by

$$\mathbf{T}_u^{\text{TSME}} = \tilde{\Theta} \tilde{\Lambda} \tilde{\Theta}_y^H (\tilde{\Theta}_y \tilde{\Lambda} \tilde{\Theta}_y^H + \tilde{\mathbf{S}}_{nn})^{-1}. \quad (28)$$

If the weighted resolvent forcing were truly uncorrelated, such that $\tilde{\Sigma}^2 = \tilde{\Lambda}$ and $\tilde{\Psi} = \tilde{\Theta}$, (21) and (28) would be identical. As such, the TRME estimation method can be used to evaluate the validity of the assumed forcing statistics and thereby the ability of the resolvent modes to approximate SPOD modes for flow reconstruction.

III. METHODS AND APPROACH

As a prototypical real-time application of the estimation techniques discussed in Sec. II, we reconstruct turbulent velocity fluctuations in a minimal channel by continuously incorporating streaming measurements in a manner akin to time-frequency analysis. Within this streaming context, we refer to the techniques in Secs. IIC 1, IIC 2, and IIC 3 as streaming TRME (STRME), streaming ORBE (SORBE), and streaming TSME (STSME), respectively. In what follows, we detail our methods and highlight how we enable and evaluate efficient streaming flow reconstructions.

A. Numerical simulation details

We apply the flow reconstruction framework to reconstruct velocity fluctuations from a DNS [53] of turbulent channel flow in a minimal flow unit [54] at $\text{Re}_\tau \approx 186$, as discussed previously by Bae *et al.* [55]. As depicted in Fig. 2, the simulation data exist on a staggered grid with $(N_x, N_y, N_z) = (32, 129, 32)$ cells in each direction, excluding ghost cells. The grid resolution is uniform in x and z and follows a hyperbolic tangent distribution in y . The governing equations are spatially discretized using second-order central finite differences and time-marched using an explicit third-order Runge-Kutta scheme in the simulation. Table I summarizes relevant spatiotemporal parameters associated with the data we use from the DNS. While ν and h are fixed, u_τ is estimated over a long time series and assumed to be constant. In the present work, we subsample the DNS time steps such

TABLE II. Measurement configurations for the testing period, including the number of wall-parallel planes (N_{plane}) and the wall-normal locations (y^+) in viscous units.

Case	N_{plane}	y^+
A	1	14.7
B	2	14.7 (both sides)
C	4	14.7, 56.4 (both sides)
D	6	14.7, 56.4, 114 (both sides)
E	7	14.7, 56.4, 114 (both sides), 186 (centerline)

that $\Delta t^+ > \Delta t_{\text{DNS}}^+$ to efficiently process snapshots without significantly affecting the resolution of relevant temporal frequencies.

B. Measurement details

We reconstruct velocity fluctuations using triplet space measurements that may evolve over time in a manner similar to time-frequency analysis. Measurements and reconstructions are restricted in the present study to the smallest $(N_{k_x}, N_{k_z}) = (3, 5)$ nonnegative streamwise and spanwise wave numbers and their complex conjugates. While this implies that we truncate more than 95% of the Fourier modes corresponding to the spatial discretization scheme, these modes still contain a dominant portion of the response energy. The temporal Fourier modes are computed by discretizing a single eddy turnover time (i.e., $\Delta T = 1$) and consist of $N_\omega = 176$ nonnegative frequencies and their complex conjugates. These selections allow us to capture convection velocities more than $8\times$ larger than the centerline velocity, although most energetic motions propagate significantly slower. Considerable energy is contained by modes with periods longer than $\Delta T = 1$, but the present configuration represents a balance of capturing energetic Fourier modes and retaining efficient reconstructions. The zero values in each triplet are numerically set to $O(10^{-10})$, and we verify that the transfer functions are invariant to this threshold.

We use simple heuristic arguments to inform sensor placement. As summarized in Table II, we consider a spectrum of measurement configurations ranging from relatively cheap (case A) to relatively expensive (case E). Case A represents a single sensor plane coinciding with the peak streamwise root-mean-square (rms) velocity fluctuations, u_{rms} . It is unique in that it evaluates the efficacy of sensors localized to one half of the channel. Case B is similar but symmetrically extends the configuration about the centerline. Case C extends the configuration of Case B to also consider measurements near the peaks of v_{rms} and w_{rms} . Case D adds sensor planes between these peaks and the centerline to better capture the outer structure of the fluctuations. Case E adds the centerline to evaluate whether it significantly improves the quality of the reconstructions.

Using these restricted numerical data, we simulate the process of obtaining experimentally feasible measurements. An analogous experimental configuration involves using standard multiplane PIV techniques [56,57] to obtain velocity measurements in wall-parallel planes. Figure 3 shows a diagram of this configuration in the context of real-time flow reconstructions. This configuration contrasts slightly with standard $x - y$ planar PIV setups, and inhomogeneities associated with bounding walls may need to be modeled into the noise CSD. Experimental reconstructions would require windowed spatial Fourier transforms in x and z since periodicity cannot be guaranteed, but this complication may also be absorbed directly into the noise vector. When designing PIV measurements, the relevant streamwise and spanwise scales of large-scale motions must be known (or modeled) to accurately capture their Fourier space representations. However, since the measurements consistently update the phase information of the estimated fluctuations, this restriction may be mitigated by absorbing unresolved large-scale motions into the streamwise-constant modes at each time step.

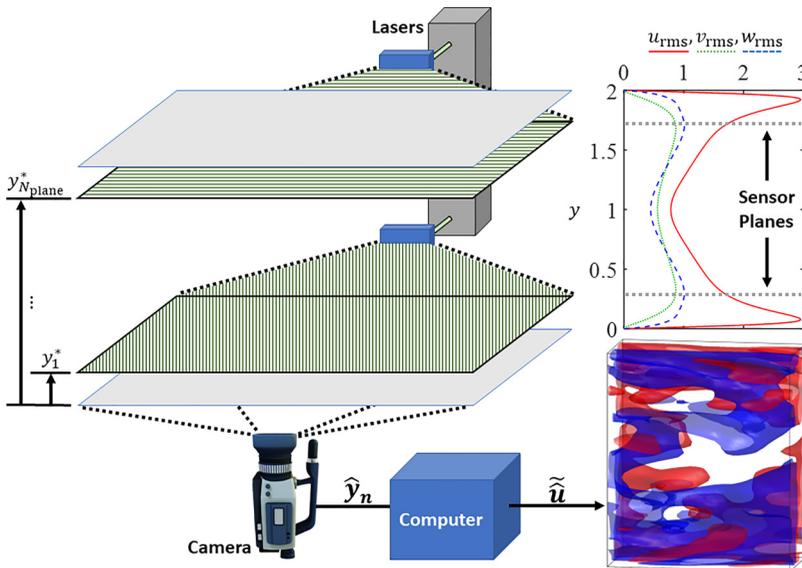


FIG. 3. Analogous multiplane PIV experimental configuration that is amenable to real-time reconstructions of velocity fluctuations under the present framework. The lines along each laser sheet represent orthogonal polarizations that enable simultaneous measurements, and more than one camera may be needed to fully capture all measurement planes [56,57]. The wall-normal sensor locations with superscripts $(\cdot)^*$ are selected from Table II and represent testing period measurements.

C. Training and testing periods

The reconstruction techniques we use require *a priori* knowledge of various linear models and flow statistics to form linear estimators. To address this, we partition the DNS data into training and testing periods of equal size such that $\Delta T' = \Delta T^* = 80$, where $(\cdot)'$ and $(\cdot)^*$ represent training and testing period quantities, respectively. As summarized in Fig. 4, we form each estimator using flow statistics from the training period and evaluate the performance of streaming flow reconstructions during the subsequent testing period. While u_τ is assumed to be known prior to training, the error in u_τ estimated over the training period with respect to its assumed constant value is 0.172%. Further, all present flow estimators and their reconstruction accuracies are relatively insensitive to $\Delta T'$ and ΔT^* , respectively.

STRME, SORBE, and STSME require various degrees of training data to form the corresponding estimators. STRME requires only a mean profile to compute the resolvent operator, which may be efficiently estimated using ergodicity or modeled using scaling laws. As discussed in Sec. III C 1, we introduce an efficient and accurate blockwise inversion technique to compute the resolvent operator for STRME and SORBE. SORBE requires more detailed (second-order) statistics, $S'_{yy,n}$, for an arbitrary number of auxiliary sensor planes. By limiting the number of sensor planes considered, SORBE provides an efficient means of estimating global flow statistics from limited training data. At the high end, STSME requires full-field second-order statistics, S_{uu} , to compute SPOD modes and energies. In this limit of complete statistical information, both (untruncated) STSME and SORBE approach the generalized Wiener filter transfer function, $S_{uy}S_{yy,n}^{-1}$.

For SORBE and STSME, we assume 50% overlap between adjacent $\Delta T = 1$ windows when computing the relevant CSDs. This gives $N_{\text{real}} = 158$ realizations of each measurement set over the training period. We intentionally design these experiments such that $N_{\text{real}} < N_u$ to evaluate the efficacy of our statistical methods under typical conditions for turbulence estimation problems. We do not use windowing in computing the CSDs to preserve the relationships associated with the forcing and response (CSDs) without considering a correction term [41,42,58].

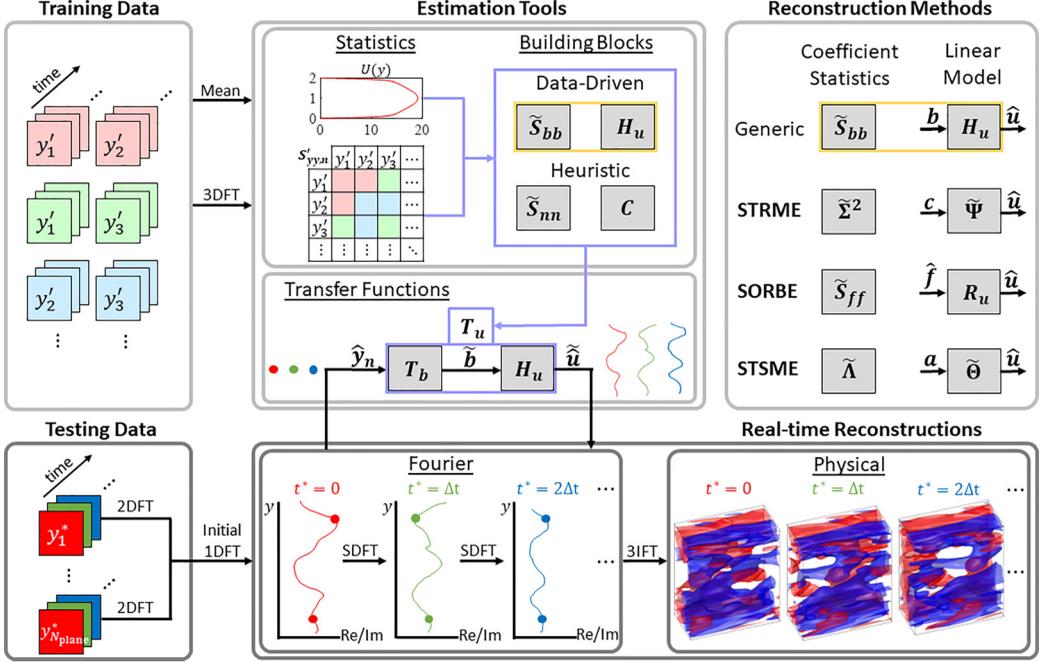


FIG. 4. Diagram depicting the methods used to estimate linear transfer functions for each reconstruction method during training (top) and to reconstruct velocity fluctuations in real time from sparse measurements during testing (bottom). Discrete Fourier transforms (DFTs) and inverse discrete Fourier transforms (IFTs) are labeled by their dimensionality. SDFT refers to a (recursive) sliding temporal DFT applied every time step, as discussed in Sec. III C 2.

For each measurement configuration in Table II, we consider estimators of varying fidelity based on auxiliary measurements (SORBE) and modal truncations (STRME and STSME). We summarize the fidelities of these estimators in Table III.

As depicted in Fig. 4, during training we consider various pairs of wall-normal coordinates, (y'_i, y'_j) , associated with auxiliary measurements that are distinct from the fixed measurements used for reconstruction. We may compute each element of $S'_{yy,n}$ (for SORBE) or S_{uu} (for STSME)

TABLE III. Subcases considered for each reconstruction method, characterized by the extent of auxiliary sensors (SORBE) and the rank of modal truncations (STRME and STSME). For each case, the subcases are listed in order of increasing fidelity. For SORBE, the y^+ inequalities apply to both sides of the channel.

Method	Parameter	Values
SORBE	y'	y^* only
		y^* and $y^+ \leq 14.7$
		y^* and $y^+ \leq 56.4$
		y^* and $y^+ \leq 114$
		all y
STRME	N_{mode}	2, 8, 40, 158, 388
STSME	N_{mode}	2, 8, 40, 158

in parallel for all \mathbf{k} using two simultaneous planar (PIV) sensors. In experiments, this implies that the length of the training period is dictated by the number of sensor pairs, N_{yy} , required to sufficiently resolve the wall-normal structure of the second-order statistics. However, under the assumption of statistical stationarity, we simultaneously collect all elements of the relevant CSDs during the training period for the present (numerical) investigations. In effect, this parallelizes the process of simulating two-plane PIV measurements for each sensor pair, which would be conducted sequentially in experiments absent extra sensors. We emphasize that N_{yy} does not necessarily scale as N_y^2 since an underlying model with sufficient fidelity can often capture the structure of the statistics with relatively few spatial measurements [3,42].

During the subsequent testing period, we use the estimators from training to reconstruct the flow with minimal time delay, as is required in real-time applications. One advantage of computing the transfer functions during training is that the flow at each triplet is efficiently reconstructed with a single application of the transfer function (via matrix multiplication). We further enable fast reconstructions by applying this operation simultaneously for all selected triplets. Since no sensitive computational techniques (e.g., inverses) are required after training, we truncate the transfer functions to single precision to improve memory usage and temporal efficiency without significantly impacting reconstruction quality.

Since an infinite time series is not available, we estimate the temporal Fourier coefficients using a relatively small sliding window ($\Delta T = 1$). We accommodate real-time estimation problems by continuously updating the temporal Fourier coefficients using the incoming data stream. As discussed in Sec. III C 2, we enable efficient streaming reconstructions by recursively updating the coefficients using the SDFT as new measurements arrive. The errors associated with the finite temporal window are implicitly accounted for in $\hat{\mathbf{n}}$, highlighting that the assumption that $\check{\mathbf{S}}_{nn} = \epsilon \mathbf{I}$ [41,43] is indeed a simplified model in the present investigation. Appendix B shows that the short-time ($\Delta T = 1$) and long-time ($\Delta T' = 80$) Fourier amplitude spectra are consistent with each other for a selected wave-number pair. For all reconstructions, we assume that $\epsilon = 10^{-8}$ to approach the zero-noise limit.

1. Blockwise inversion

In the present work, we invoke the resolvent operator when computing the STRME and SORBE estimators during training. More generally, computing the resolvent operator and forming its modal decomposition are central challenges in efficiently and accurately representing fluctuation dynamics. Efficient randomized [59] and matrix-free [60] techniques for forming these equation-based elements have been studied previously. Here we introduce the exact blockwise representation of the linear dynamics to enable efficient and accurate computation of the channel flow resolvent operator by means of blockwise inversion. This methodology takes advantage of the simple blockwise structure associated with 1D resolvent analysis and expresses the inverse of the full linear dynamics in terms of its smaller subblock constituents.

Conventionally, the resolvent operator in (11) is obtained by directly inverting \mathbf{L} and requires a large matrix inverse. However, the form of \mathbf{L} in (4) lends itself to blockwise inversion. We begin by blockwise inverting the basic subblock of \mathbf{L} to give the basic resolvent operator,

$$\mathbf{R}_u^B = \mathbf{L}_B^{-1} = \begin{bmatrix} \mathbf{L}_c^{-1} & -\mathbf{L}_c^{-1}(\partial_y U)\mathbf{L}_e^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_e^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{L}_c^{-1} \end{bmatrix}, \quad (29)$$

which is defined here between forcing and response quantities that are aligned in x and z for convenience (see Sec. II B). The simple inverse defining the basic resolvent neglects the contribution of the pressure gradient term and that only the solenoidal forcing and response are active in the input-output formulation for velocity fluctuations [42,51]. Now, by blockwise inverting \mathbf{L} and applying the definition of the resolvent in (11), we express the resolvent operator in terms of the

basic resolvent and the continuity subblock as

$$\mathbf{R}_u = \mathbf{S}^H [\mathbf{R}_u^B - \mathbf{R}_u^B \hat{\mathbf{V}} (\hat{\mathbf{V}}^T \mathbf{R}_u^B \hat{\mathbf{V}})^{-1} \hat{\mathbf{V}}^T \mathbf{R}_u^B] \mathbf{S}. \quad (30)$$

This equation may be used to rewrite the resolvent formulation in (11) as

$$\hat{\mathbf{u}} = \mathbf{S}^H [\mathbf{I} - \mathbf{R}_u^B \hat{\mathbf{V}} (\hat{\mathbf{V}}^T \mathbf{R}_u^B \hat{\mathbf{V}})^{-1} \hat{\mathbf{V}}^T] \mathbf{R}_u^B \mathbf{S} \hat{\mathbf{f}} \quad (31)$$

$$= \mathbf{S}^H \mathbf{R}_u^B [\mathbf{I} - \hat{\mathbf{V}} (\hat{\mathbf{V}}^T \mathbf{R}_u^B \hat{\mathbf{V}})^{-1} \hat{\mathbf{V}}^T \mathbf{R}_u^B] \mathbf{S} \hat{\mathbf{f}}, \quad (32)$$

where the terms in square brackets represent projection matrices since they satisfy the property that $\mathbf{P}^2 = \mathbf{P}$. These terms share a similar form with the Leray projection typically used to remove \hat{p} , except for the inclusion of the basic resolvent. Since this formulation also serves to enforce continuity and eliminate pressure fluctuations [61], we refer to it as a ‘‘modified projection.’’ The validity of applying the projection terms both before and after applying the basic resolvent highlights the correspondence between the solenoidal forcing and response for velocity fluctuations. Similar expressions are readily derived for the continuous linear system, (3), by expressing $\hat{\mathbf{u}}$ in terms of $\hat{\mathbf{f}}$ and $\nabla \hat{p}$ and enforcing continuity to eliminate the pressure term.

There are numerous advantages to computing the resolvent operator using blockwise inversion as opposed to directly inverting \mathbf{L} . First, the matrix inversions (of \mathbf{L}_c , \mathbf{L}_e , and $\hat{\mathbf{V}}^T \mathbf{R}_u^B \hat{\mathbf{V}}$) required using blockwise inversion are roughly 4×4 smaller than that of \mathbf{L} . Further, if flow variables are collocated (such that $\mathbf{L}_c = \mathbf{L}_e$), computing the basic and full resolvent operators requires only one and two small inverses, respectively. The advantage of these smaller matrix inverses is that the smaller subblocks have smaller condition numbers than the full linear system. Therefore, blockwise inversion is typically more computationally efficient and accurate than direct inversion of the entire system. Further, the form of the correction term in (30) reveals a direct relationship between the basic resolvent operator, where only the terms in the momentum equations that are linear in $\hat{\mathbf{u}}$ are inverted, and the full resolvent operator, which includes the continuity constraint and eliminates the pressure fluctuations. Therefore, the present blockwise formulation of the resolvent operator allows for a physical characterization of the influence of various subblocks (representing different facets) of the linear dynamics. We apply this technique to the resolvent operator modified to include an eddy-viscosity model [17,43] in Appendix A, but we do not employ that model in the present computations.

2. Recursive sliding discrete Fourier transform

During the testing period, we employ the sliding discrete Fourier transform (SDFT) [46,62,63] to recursively update the temporal Fourier coefficients of the measurements at every time step, t_i . We thereby enable efficient updates of our triplet space measurements as new data arrive. We compute discrete Fourier coefficients using a sliding temporal window of fixed length ($\Delta T = 1$) that is discretized into N_t time steps. Given a (potentially complex) scalar function, $\phi(t)$, we assume that initial discrete Fourier coefficients, $\hat{\phi}_{N_t}(\omega_m)$ for $m = 1, \dots, N_t$, are computed using time steps $i = 1, \dots, N_t$. The discrete frequencies are given by $\omega_m = \Delta\omega(m - 1)$ for $m \leq \lceil N_t/2 \rceil$ and $\omega_m = \Delta\omega(m - 1 - N_t)$ for $m > \lceil N_t/2 \rceil$, where $\Delta\omega = 2\pi/N_t \Delta t$ is the frequency resolution. Using the standard SDFT, the m th discrete Fourier coefficient at a time step $i > N_t$ is computed recursively as

$$\hat{\phi}_i(\omega_m) = \underbrace{\exp(i\omega_m \Delta t)}_{\text{Phase Update}} [\underbrace{\hat{\phi}_{i-1}(\omega_m)}_{\text{Previous Value}} + \underbrace{(\phi(t_i) - \phi(t_{i-N_t}))}_{\text{Streaming Info}})]. \quad (33)$$

This process may be continuously applied as new time steps become available, thereby making it amenable to real-time computation of Fourier coefficients over a sliding temporal window.

A primary advantage of the SDFT is that a complete Fourier transform need only be computed over the initial window. Stepping forward in time, the coefficients at time step i are redundant with respect to those at time step $i - 1$ except for the difference between the incoming measurement (at i)

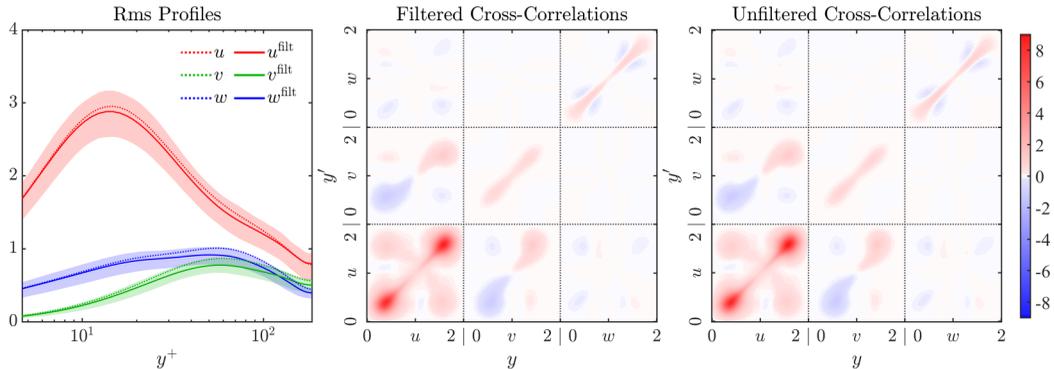


FIG. 5. Comparisons between the filtered and unfiltered flow statistics for the testing period. (Left) Comparison between the filtered rms profiles during the testing period (solid curves: square root of mean variance, shading: 20th–80th percentiles) and the long-time unfiltered rms profiles (dotted curves). For the testing period, the mean filtered cross-correlations (middle) are shown alongside their unfiltered counterparts (right).

and the outgoing measurement (at $i - N_t$). As such, instead of reusing the `fft` algorithm, for all time steps $i > N_t$ we may use the SDFT to express the Fourier coefficients recursively in terms of (i) the previous Fourier coefficients, (ii) the aforementioned difference, and (iii) a complex multiplicative correction factor. Correspondingly, the number of required operations scales as $O(N_t)$ as opposed to $O(N_t \log N_t)$ for the standard `fft`. Further, since the frequency bins may be updated independently of one another, this complexity is reduced to $O(N_\omega)$ when considering a limited subset of Fourier modes (i.e., if $N_\omega < N_t$).

The standard SDFT algorithm suffers from marginal stability, moderate accuracy, and time-accumulating error (such that, eventually, the `fft` algorithm would need to be called to refresh the error). The more recent modulated SDFT (mSDFT) [62] and observer-based SDFT (oSDFT) [63] have more desirable stability, accuracy, and robustness to noise compared to the standard SDFT and other variants [46]. However, we retain the standard SDFT due to its simple implementation and since it remains accurate and stable over sufficiently large intervals for our purposes (see Appendix B). The mSDFT, oSDFT, or some other robust method may be considered to mitigate the propagation of noise over longer time intervals.

D. Error metrics

A primary objective of our methodology is to produce efficient flow reconstructions that accurately capture the dynamics throughout the channel. Here we define error metrics that reflect reconstruction accuracy at various levels of detail. These metrics are advantageous in that they enable comparisons of spatially integrated errors and their wall-normal distributions across different estimators. Whereas previous investigations [17,43] have considered similar metrics, they focus more on errors at specific estimation planes and wave numbers.

Noting that we use only a subset of roughly 5% of the spatial Fourier modes in the channel, we begin by characterizing how well the filtered dynamics associated with these modes capture unfiltered (i.e., true) fluctuations. As shown in Fig. 5, the filtered dynamics excellently capture the statistical features of the unfiltered dynamics. On average, the filtered fluctuations during the testing period contain 91.1% of the turbulence kinetic energy (TKE) of the long-time unfiltered fluctuations. Further, the (wall-normal) cross-correlations associated with filtered and unfiltered dynamics have an excellent agreement over the testing period. These results justify our compressed triplet space flow representation and motivate the definition of the following error metrics.

The error in TKE coarsely characterizes the accuracy of the reconstructed turbulence intensity. Correspondingly, we evaluate the turbulence intensity error using the square root of the relative TKE

error, given by

$$\epsilon_{\text{TKE}}(t) = \left(\frac{\int_0^2 |\mathbf{u}_{\text{rms}}(y, t)|^2 dy - \int_0^2 |\tilde{\mathbf{u}}_{\text{rms}}(y, t)|^2 dy}{\int_0^2 |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad (34)$$

where rms fluctuations are computed using the sliding temporal window ($\Delta T = 1$). This sliding window allows all error metrics to take into account time-varying flow statistics in streaming applications, thereby enabling quantification of temporal uncertainties.

Beyond reconstructing flow statistics, we are primarily interested in the accuracy of the reconstructed dynamics. We consider two error profiles, $\epsilon_{\text{filt}}(y, t)$ and $\epsilon_{\text{full}}(y, t)$, that capture the wall-normal distributions of the reconstruction errors with respect to the filtered and unfiltered snapshots, respectively. These error profiles are given by

$$\epsilon_{\text{filt}}(y, t) = \left(\frac{1}{N_x N_z} \|\mathbf{u}_{\text{filt}}(\mathbf{x}, t) - \tilde{\mathbf{u}}(\mathbf{x}, t)\|_{x,z}^2 \right)^{1/2}, \quad (35)$$

$$\epsilon_{\text{full}}(y, t) = \left(\frac{1}{N_x N_z} \|\mathbf{u}(\mathbf{x}, t) - \tilde{\mathbf{u}}(\mathbf{x}, t)\|_{x,z}^2 \right)^{1/2}, \quad (36)$$

where \mathbf{u}_{filt} represents the filtered snapshots and $\|\cdot\|_{x,z}$ represents the L^2 norm computed over the streamwise and spanwise directions.

We also evaluate time-varying full-field errors by integrating the error variance profiles over the wall-normal direction. The corresponding integrated error metrics, which are given by

$$\epsilon_{\text{filt}}(t) = \left(\frac{\int_0^2 |\epsilon_{\text{filt}}(y, t)|^2 dy}{\int_0^2 |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad \epsilon_{\text{full}}(t) = \left(\frac{\int_0^2 |\epsilon_{\text{full}}(y, t)|^2 dy}{\int_0^2 |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad (37)$$

evaluate (under the square root) the error energies with respect to the TKE. Importantly, we use the same (unfiltered) TKE in both $\epsilon_{\text{filt}}(t)$ and $\epsilon_{\text{full}}(t)$ to simplify the visual comparisons between both metrics, but using the filtered TKE for $\epsilon_{\text{filt}}(t)$ also provides a useful metric.

Since we primarily reconstruct the flow using a small number of sparsely sampled measurements, it may not be reasonable to expect highly accurate reconstructions far from the measurement planes. This premise is especially true for flow structures that are localized in the wall-normal direction. To address this issue, we introduce truncated error metrics, $\epsilon_{\text{TKE}}^*(t)$, $\epsilon_{\text{filt}}^*(t)$, and $\epsilon_{\text{full}}^*(t)$, that evaluate the errors in regions localized about the measurement planes. These metrics are given by

$$\epsilon_{\text{TKE}}^*(t) = \left(\frac{\int_{y_{\text{local}}} |\mathbf{u}_{\text{rms}}(y, t)|^2 dy - \int_{y_{\text{local}}} |\tilde{\mathbf{u}}_{\text{rms}}(y, t)|^2 dy}{\int_{y_{\text{local}}} |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad (38)$$

$$\epsilon_{\text{filt}}^*(t) = \left(\frac{\int_{y_{\text{local}}} |\epsilon_{\text{filt}}(y, t)|^2 dy}{\int_{y_{\text{local}}} |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad \epsilon_{\text{full}}^*(t) = \left(\frac{\int_{y_{\text{local}}} |\epsilon_{\text{full}}(y, t)|^2 dy}{\int_{y_{\text{local}}} |\mathbf{u}_{\text{rms}}(y, t)|^2 dy} \right)^{1/2}, \quad (39)$$

where y_{local} includes all cells within N_{local} cells from each measurement plane. These definitions allow us to balance reconstruction accuracy with the fraction of the channel captured. Moreover, since Δy_j is smaller close to the walls and larger close to the centerline, these metrics provide a simple means of adapting the extent over which the errors are integrated based on the expected scale of the corresponding flow structures.

One unique advantage of the ϵ_{TKE} metrics in (34) and (38) is that they have the potential to be evaluated in experimental settings. By contrast, the ϵ_{filt} and ϵ_{full} metrics in (37) and (39) require access to instantaneous 3D flow fields, which are typically not available in experiments. However, one disadvantage of the ϵ_{TKE} metrics in isolation is that their values do not reflect the wall-normal

distribution of TKE. Nevertheless, the TKE-based error metrics remain useful when supplemented with the dynamics-based error metrics since, together, they encode information regarding the spatially integrated correlations between the reconstructions and the DNS.

IV. RESULTS

In this section we evaluate the accuracy and efficiency of flow reconstructions using STRME, SORBE, and STSME.

A. Equation-based estimation: STRME

Using no training data except a mean profile, STRME forms a simple, informative framework for flow reconstruction. However, its accuracy relies on an unrealistic model of an uncorrelated forcing. Figures 6 and 7 summarize the accuracy of the STRME estimators for each modal truncation. Due to the unrealistic uncorrelated forcing model, the global errors are relatively large for all modal truncations.

For $N_{\text{plane}} \geq 6$, the $N_{\text{mode}} = 8$ ROM produces the most accurate full-field reconstructions in a relatively economical fashion. However, for $N_{\text{plane}} \leq 6$ the $N_{\text{mode}} = 2$ ROM outperforms all higher-order truncations considered. These results suggest that the uncorrelated forcing assumption is most valid for the gain-dominant resolvent modes, for which the gains are likely to contribute significantly to the mode weight statistics. Since the validity of the forcing assumption underlies correspondence between resolvent modes and SPOD modes, our results are consistent with previous observations [10] that this correspondence is most appropriate for dominant resolvent modes. The low-rank truncations for STRME and STSME are qualitatively similar in the present investigation, but the STSME estimators are more accurate as they are derived from higher-fidelity training data (see Sec. IV C).

The higher-order truncations are nearly identical visually (e.g., for $N_{\text{mode}} = 158$ and 388) since the corresponding gains are too small to impact the estimator. Especially for truncations with $N_{\text{mode}} > N_{\text{plane}}$, we expect that improved models of the mode weight statistics are likely to improve estimator performance. Further, in the limit of including all $N_u = 388$ modes, STRME represents a form of SORBE under the uncorrelated forcing assumption. One natural, data-driven extension of STRME involves estimating the mode weight statistics using limited training data (as is done for the forcing in SORBE). However, as shown in Appendix C, the forcing-based SORBE estimator empirically outperforms the mode weight-based estimator inspired by STRME in the present investigation. In the context of the open question of optimal coefficients for flow reconstruction, this justifies our use of the forcing over the mode weights in the SORBE formulation.

Despite relatively poor full-field accuracy, the STRME estimators are more accurate in the vicinity of the measurement planes. This effect is especially pronounced for $N_{\text{plane}} \leq 2$, although a relatively small portion of the channel is captured in these cases. The spatial profiles for $N_{\text{plane}} = 7$ confirm that reconstructions tend to be more accurate near the measurement planes in that the errors associated with higher-order modal truncations are concentrated far from the measurement planes. These results indicate that targeted reconstructions of localized flow structures are more feasible than accurate full-channel reconstructions and may require considerably fewer measurements to provide meaningful (partial) flow information.

The STRME reconstructions of u_{rms} are remarkably good considering the training data includes only a (streamwise) mean profile, but the reconstructions of v_{rms} and w_{rms} are still missing considerable energy for all STRME estimators. Consistent with previous findings [32], these results suggest that a broadband forcing assumption can be used to garner significant information about u_{rms} given only a mean profile (via the resolvent formulation). The reconstructed mean cross-correlations reinforce this notion, as the most energetic correlations captured involve the streamwise fluctuations. Further, the snapshots in Fig. 7 confirm that energetic flow structures are better captured for the streamwise fluctuations than for the wall-normal or spanwise fluctuations.

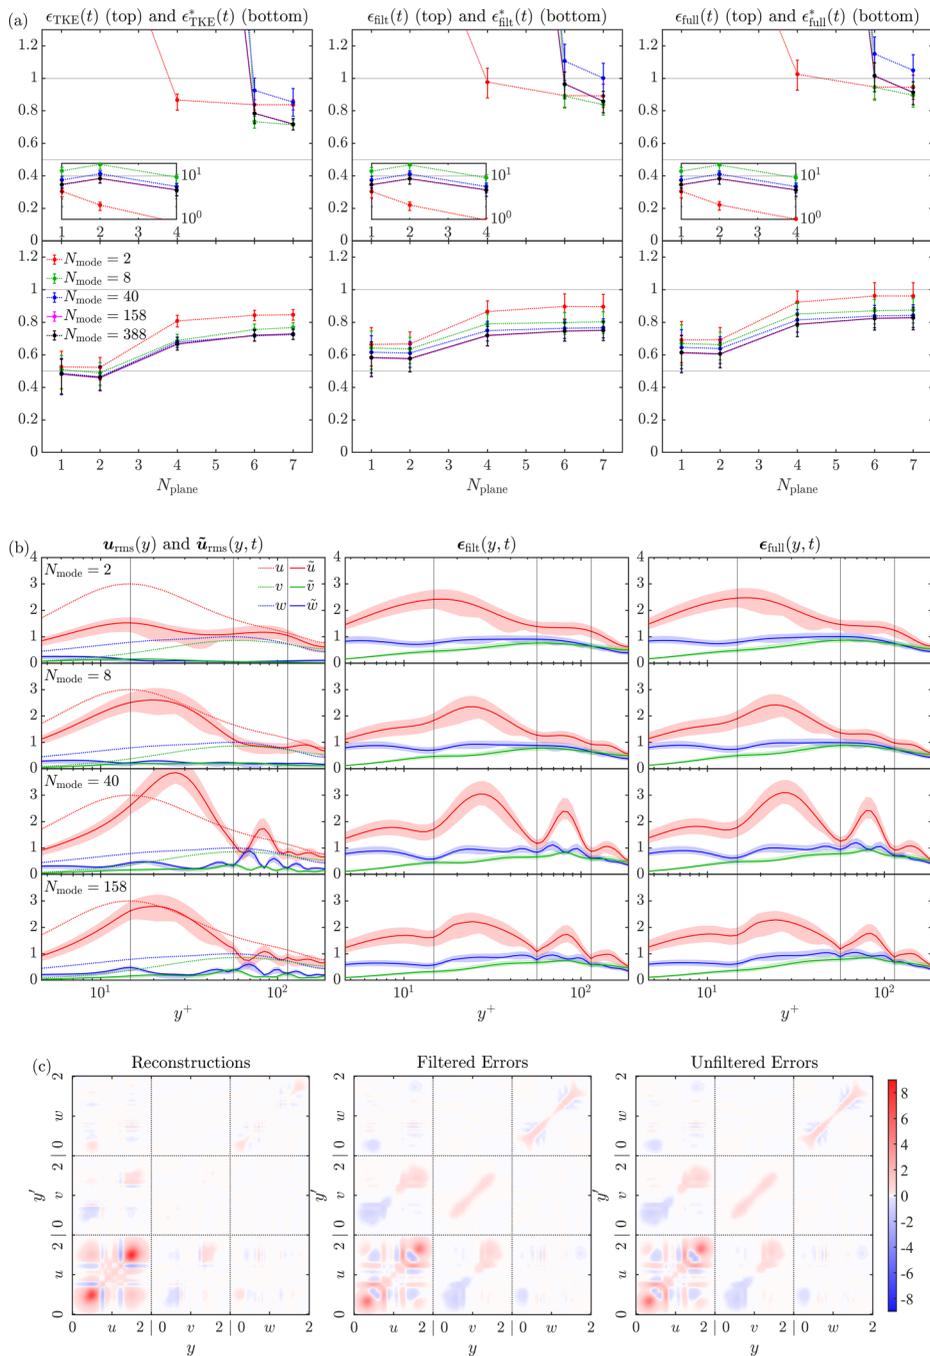


FIG. 6. Summary of reconstruction accuracies for the STRME estimators. (a) Full-field (top) and truncated (bottom) error statistics, where $N_{\text{local}} = 2$ is used for the truncated metrics. (b) RMS velocity profiles and error profiles for reconstructions using different modal truncations and $N_{\text{plane}} = 7$. The vertical lines represent the sensor locations. (c) Reconstructed cross-correlations (left) and their mean errors with respect to the filtered (middle) and unfiltered (right) cross-correlations for $N_{\text{plane}} = 7$ and $N_{\text{mode}} = 158$. The averages in (a) and (b) are computed prior to applying the square roots in the error metrics and the uncertainties extend to the 20th and 80th percentiles.

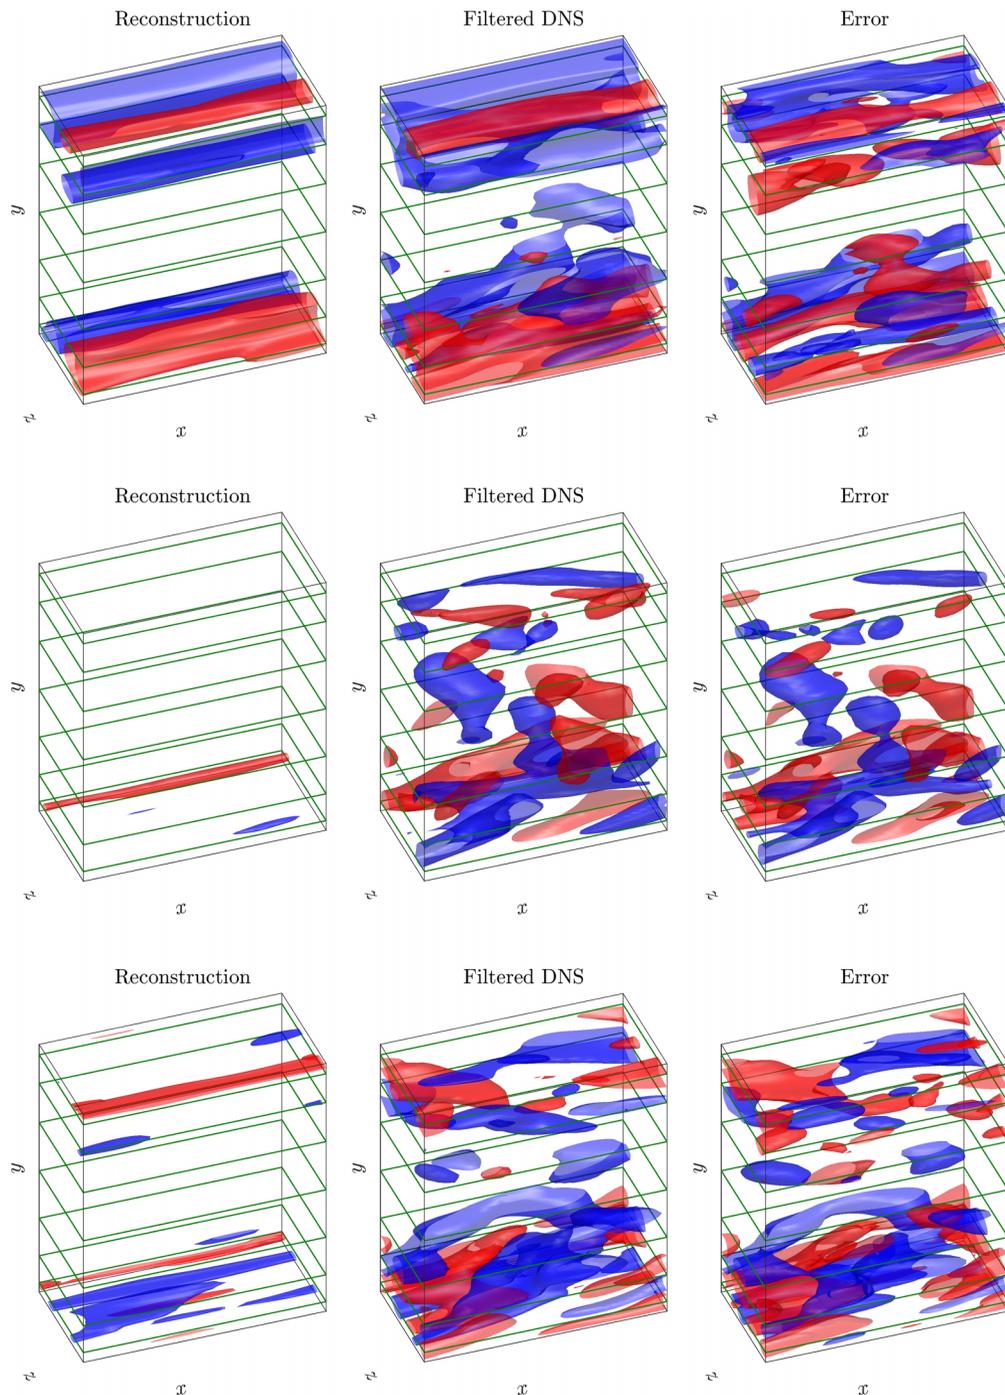


FIG. 7. Comparisons of the reconstructions, filtered DNS snapshots, and their differences for the stream-wise (top), wall-normal (middle), and spanwise (bottom) velocity fluctuations at $t = 150$. The comparisons here are shown for STRME with $N_{\text{plane}} = 7$ and $N_{\text{mode}} = 158$, for which $\epsilon_{\text{filt}}(t) = 0.708$ and $\epsilon_{\text{full}}(t) = 0.737$. Isosurfaces (+: red, -: blue) are based on the TKE of each velocity component and have magnitudes $\langle u_{\text{rms}}^2 \rangle^{1/2} = 1.58$ (top), $\langle v_{\text{rms}}^2 \rangle^{1/2} = 0.68$ (middle), and $\langle w_{\text{rms}}^2 \rangle^{1/2} = 0.77$ (bottom), where $\langle \cdot \rangle$ represents the mean over the wall-normal direction.

In summary, STRME can leverage the governing equations to produce informative localized reconstructions, especially of streamwise statistics, when no training data are available.

B. Hybrid estimation: SORBE

SORBE provides a practical approach to improve estimator performance by using spatially limited data to produce more realistic forcing statistics than the uncorrelated forcing CSD assumed in STRME. This hybrid framework provides a natural means of incorporating training measurements in conjunction with the governing equations. It allows a balance between cost and accuracy along the spectrum of training data fidelity. Figures 8 and 9 summarize the accuracy of the SORBE estimators computed using various extents of training data.

For all measurement configurations, increasing the spatial extent of training data used for the forcing estimate broadly improves reconstruction accuracy. When using relatively few measurement planes ($N_{\text{plane}} \leq 4$), the forcing fidelity has a large influence on reconstruction accuracy. By contrast, using more measurement planes ($N_{\text{plane}} \geq 6$) produces relatively smooth and more incremental (but still significant) reductions in error. These behaviors highlight two forms of the cost-accuracy tradeoff associated with augmenting the SORBE estimator using measurements. When the testing measurements are sparse, high-fidelity training data can drastically improve reconstruction accuracy. When higher-fidelity testing measurements are available, lower-fidelity training data are required to achieve similar accuracy, and high-fidelity training data further improve reconstruction accuracy.

The errors localized about each measurement plane are even lower than the full-field errors for all SORBE estimators considered. Since reconstructions are not limited by an ill-posed flow representation, the filtered errors are small in the vicinity of the measurement planes. Correspondingly, the unfiltered errors at these planes are also relatively small due to the minimal energy associated with the excluded Fourier modes. Further, as depicted in the profiles for $N_{\text{plane}} = 7$, errors tend to grow with the distance from the measurement planes. Between the measurement planes, the errors for the streamwise and spanwise fluctuations grow more rapidly than for the wall-normal fluctuations since they are (i) larger in magnitude and (ii) more localized in the wall-normal direction.

The errors in the regions between measurements are primarily attributed to the extent of the training data and the fidelity of the streaming measurements during testing. When only partial training data are used, errors grow in the gray regions (see Fig. 8) where measurements were not used to inform forcing statistics. This error source reflects that the least-squares forcing statistics informed by partial training data do not fully capture the structure of the true forcing statistics. However, even when the estimator incorporates full-field statistics from training, the errors still grow with the distance from the measurements. In the noncausal estimation setting, this error source formally reflects that the flow is poorly correlated with the measurements far from the sensor planes. Hence, while it is unreasonable to expect perfect global reconstructions from sparse measurements, this behavior can be used to inform the number of and spacings between sensors used for flow reconstruction. Moreover, as discussed more extensively in Sec. IV D, this error source is augmented by the application of the formally noncausal estimator in a streaming setting, where the flow statistics vary in time.

In the limit of full training data, SORBE produces results that converge to those produced by a generalized Wiener filter. When the generalized Wiener filter is used with $N_{\text{plane}} = 7$ (as in Case E), the mean values of $\epsilon_{\text{filt}}(t)$ and $\epsilon_{\text{full}}(t)$ are 41.5% and 52.0%, respectively. Since these errors are averaged prior to applying the square roots in their definitions, the corresponding error energies are 17.2% and 27.0% of the TKE, respectively. Further, the reconstructed mean cross-correlations are significantly more energetic than their corresponding errors for this estimator, indicating the second-order flow statistics during the testing period are well captured.

Figure 9 further reinforces that this estimator performs well throughout the channel, as the dominant global flow structures are well reconstructed for all components of the velocity fluctuations. The energetic flow structures are best captured for the wall-normal fluctuations (v) since they are

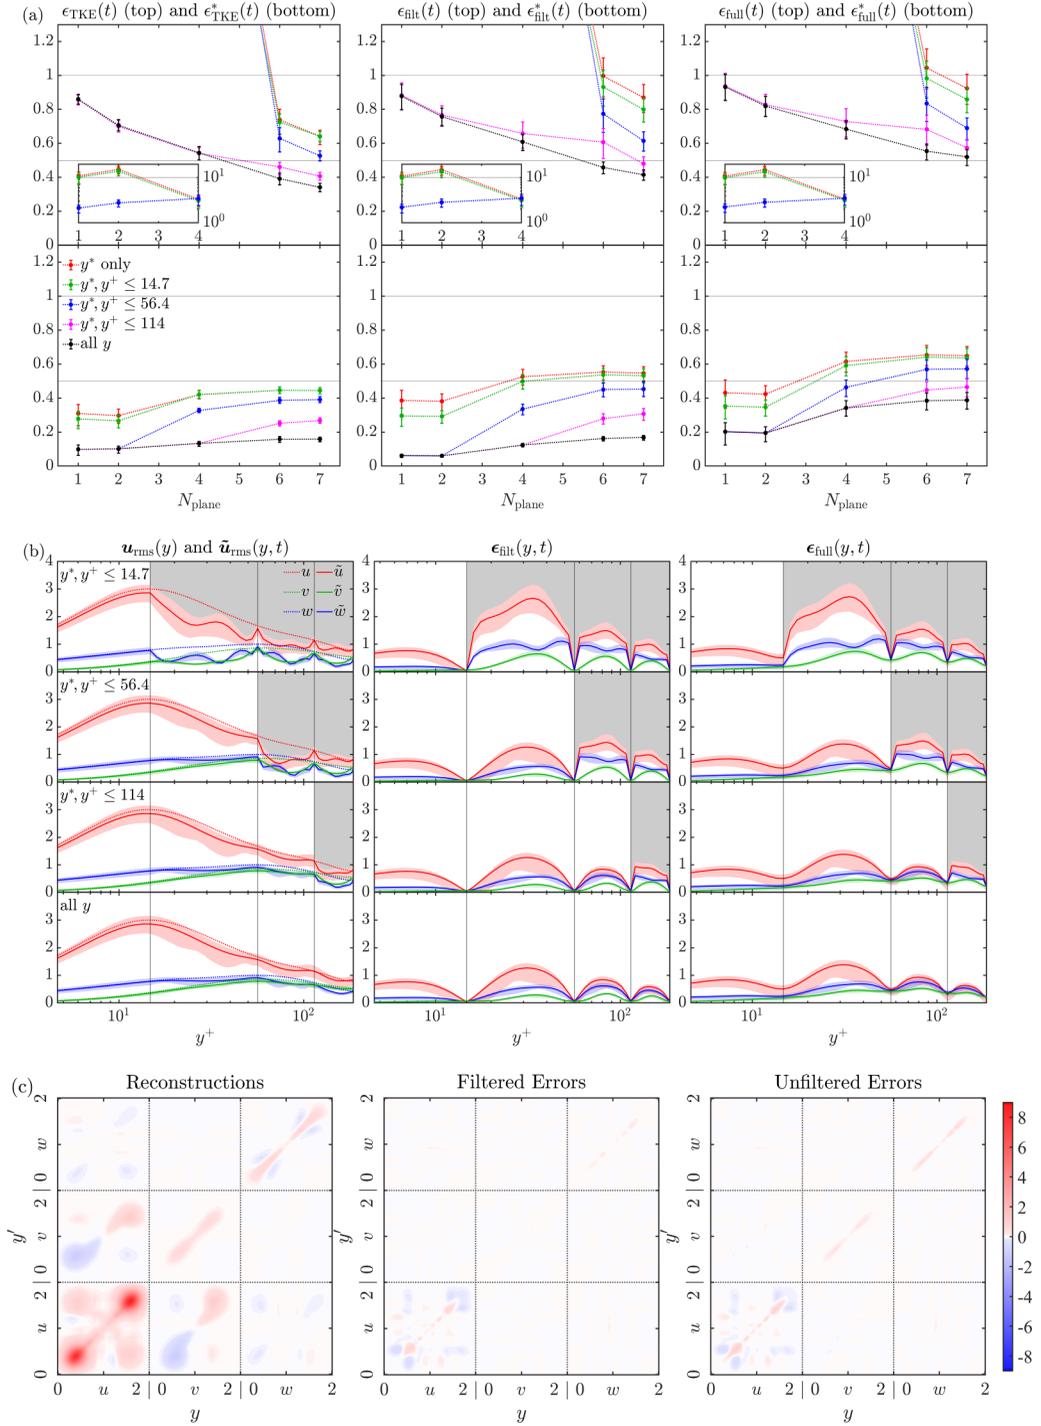


FIG. 8. Summary of reconstruction accuracies for the SORBE estimators in the style of Fig. 6. Here the spatial profiles (for $N_{\text{plane}} = 7$) are shown for various extents of auxiliary training data, where regions with gray shading are not captured during training. The cross-correlations are shown for the $N_{\text{plane}} = 7$ case with full-field training data and represent the generalized Wiener filter.

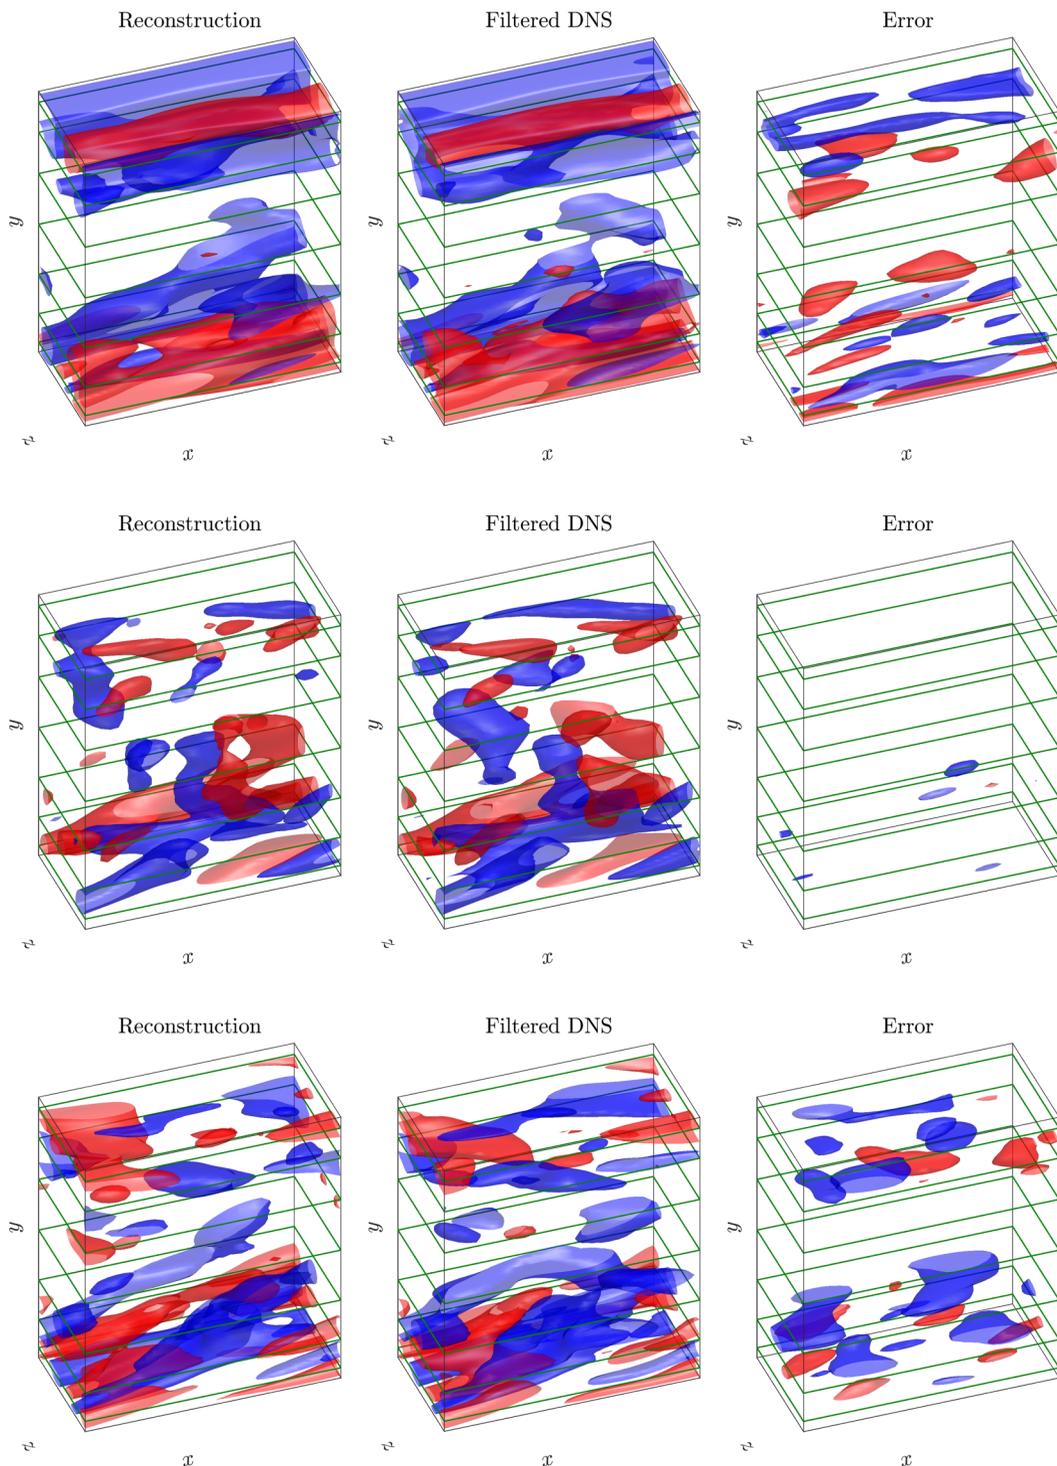


FIG. 9. Comparisons of the reconstructions, filtered DNS snapshots, and their differences in the style of Fig. 7. The comparisons here are shown for SORBE with $N_{\text{plane}} = 7$ and full-field training data (i.e., the generalized Wiener filter), for which $\epsilon_{\text{filt}}(t) = 0.399$ and $\epsilon_{\text{full}}(t) = 0.448$. Videos depicting the temporal evolutions of these comparisons are provided in the Supplemental Material [64].

less localized in y than the streamwise (u) and spanwise (w) fluctuations. Further, the reconstructions of structures of u and w tend to accumulate errors between sensor planes since the corresponding structures tend to be more localized in y . These observations are consistent with the error profiles for each component and the reductions in error when considering localized wall-normal subdomains.

In summary, the SORBE estimator provides a practical and effective means of improving reconstruction accuracy by balancing the cost-accuracy tradeoff associated with increasing training (and testing) data fidelity. In the full training data limit, a relatively small number of sensors can reproduce the dominant statistical and dynamical flow features with high accuracy.

C. Data-driven estimation: STSME

When full-field statistics from training data are available, STSME provides a data-driven method for flow reconstruction that efficiently captures energetic structures on a modal basis. Figures 10 and 11 summarize the accuracy of the STSME estimators for each modal truncation.

Similar to the STRME cases, the $N_{\text{mode}} = 2$ estimator can outperform the $N_{\text{mode}} = 8$ estimator when relatively few measurement planes are used. However, in contrast to the STRME cases, the higher-order truncations ($N_{\text{mode}} \geq 40$) produce similar, relatively accurate reconstructions for all measurement configurations. This behavior and the approximate stationarity of the estimation errors during the testing period suggest that the SPOD modes, and the flow statistics underlying their computation, are approximately converged. Hence, the coefficient statistics given by the SPOD eigenvalues accurately weight the SPOD modes, unlike in the STRME method. The performance of the $N_{\text{mode}} = 40$ estimator also suggests that the first 40 SPOD modes form an approximately well-posed basis for the flow. This observation is supported by Fig. 12, which shows that the $N_{\text{mode}} = 40$ truncation captures nearly all of the kinetic energy of the velocity fluctuations. By contrast, the $N_{\text{mode}} = 2$ and $N_{\text{mode}} = 8$ truncations capture significant but relatively incomplete portions of the kinetic energy, especially at smaller spatial scales.

Since the SPOD modes remain energetically relevant up to approximately mode 40, the lower-rank modal truncations provide ill-posed bases for representing the flow. Correspondingly, whereas the localized errors are excellent for $N_{\text{mode}} \geq 40$, they are considerably larger for the lower-order truncations. Similarly, as shown for $N_{\text{plane}} = 7$, the filtered reconstruction errors at the measurement planes are only near zero for $N_{\text{mode}} \geq 40$ and they are visibly nonzero for the lower-order cases. As observed for SORBE, even with a well-posed basis, the reconstruction errors still grow far from the measurements due to the locality of the correlations and the influence of time-varying flow statistics (see Sec. IV D).

When all SPOD modes are retained, the STSME estimator converges to the generalized Wiener filter (as did the SORBE estimator with full training data). Correspondingly, the transfer functions computed using full-field training data for SORBE and an untruncated basis for STSME produce identical error profiles and integrated errors. As N_{plane} increases, these errors decay smoothly since the flow becomes increasingly well correlated with the measurements.

Besides the generalized Wiener filter, even the lower-order truncations produce impressive reconstructions when $N_{\text{plane}} = 7$. Even though using $N_{\text{mode}} = 2$ provides a highly compressed representation of the flow (retaining only 1.2% of the modes), it produces very good reconstructions of various flow features during the testing period. For example, it produces good reconstructions of the rms fluctuations, especially for the streamwise direction. Further, the dominant features of the mean cross-correlations are very well captured during the testing period. The snapshots in Fig. 11 further support this observation. Whereas the errors are not negligible with respect to the flow structures, the reconstructions qualitatively capture most relevant flow structures for each velocity component. Hence, if leading SPOD modes can be computed using alternative means to alleviate the training data requirement, STSME has the potential to provide efficient reconstructions of energetically relevant flow structures. However, as presently formulated, the spatial extent of the training statistics is likely a prohibitive limitation to applying the STSME estimators in practice.

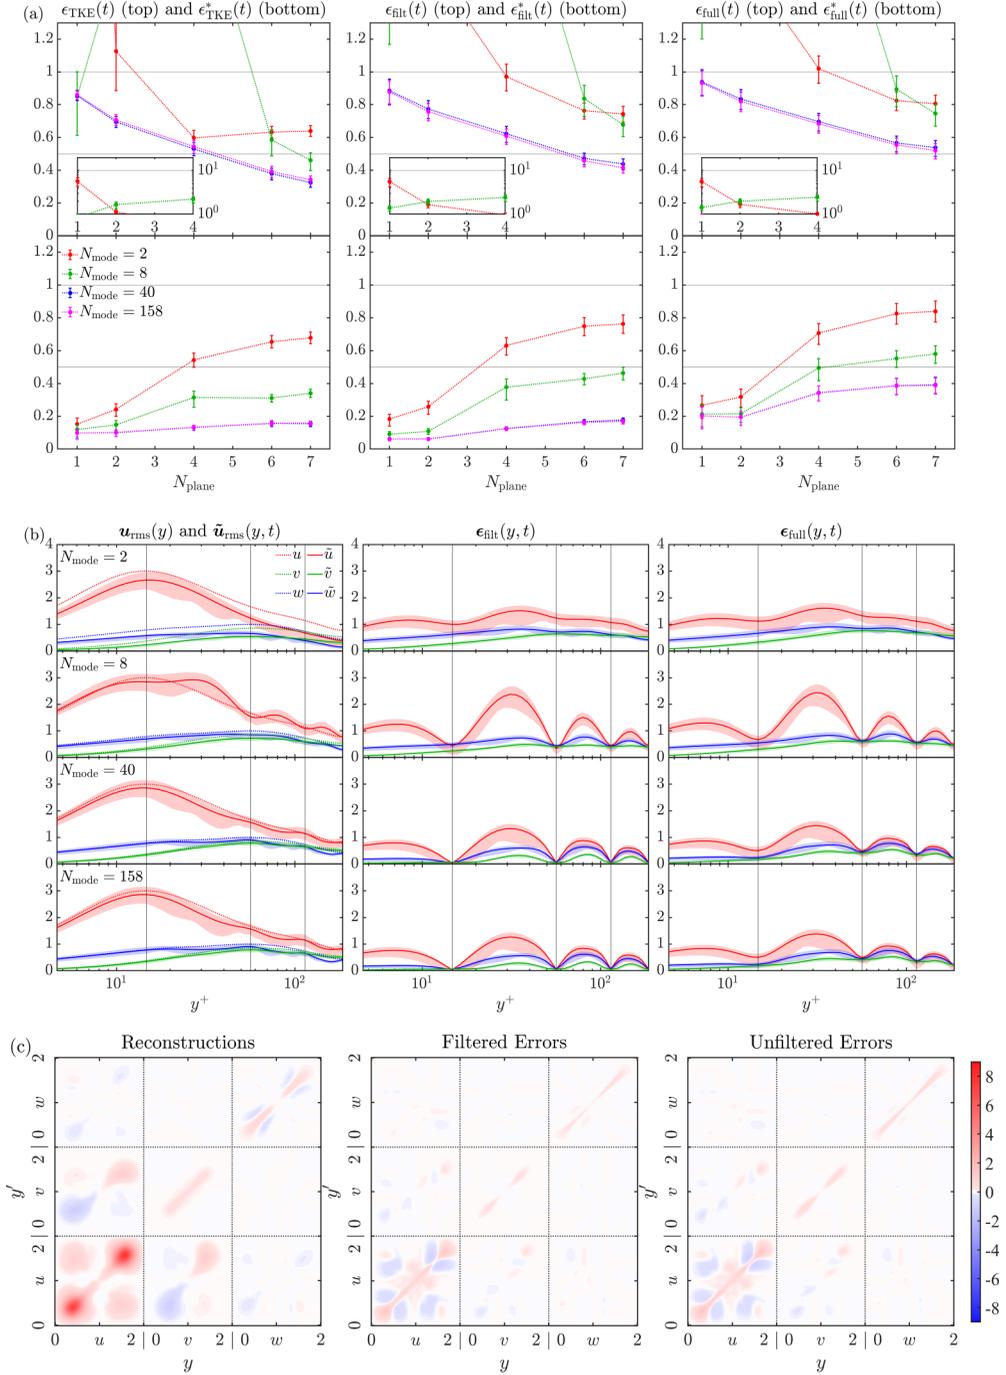


FIG. 10. Summary of reconstruction accuracies for the STSME estimators in the style of Figs. 6 and 8. Here the spatial profiles (for $N_{\text{plane}} = 7$) are shown for various modal truncations and the cross-correlations (for $N_{\text{plane}} = 7$) are shown for the case with $N_{\text{mode}} = 2$.

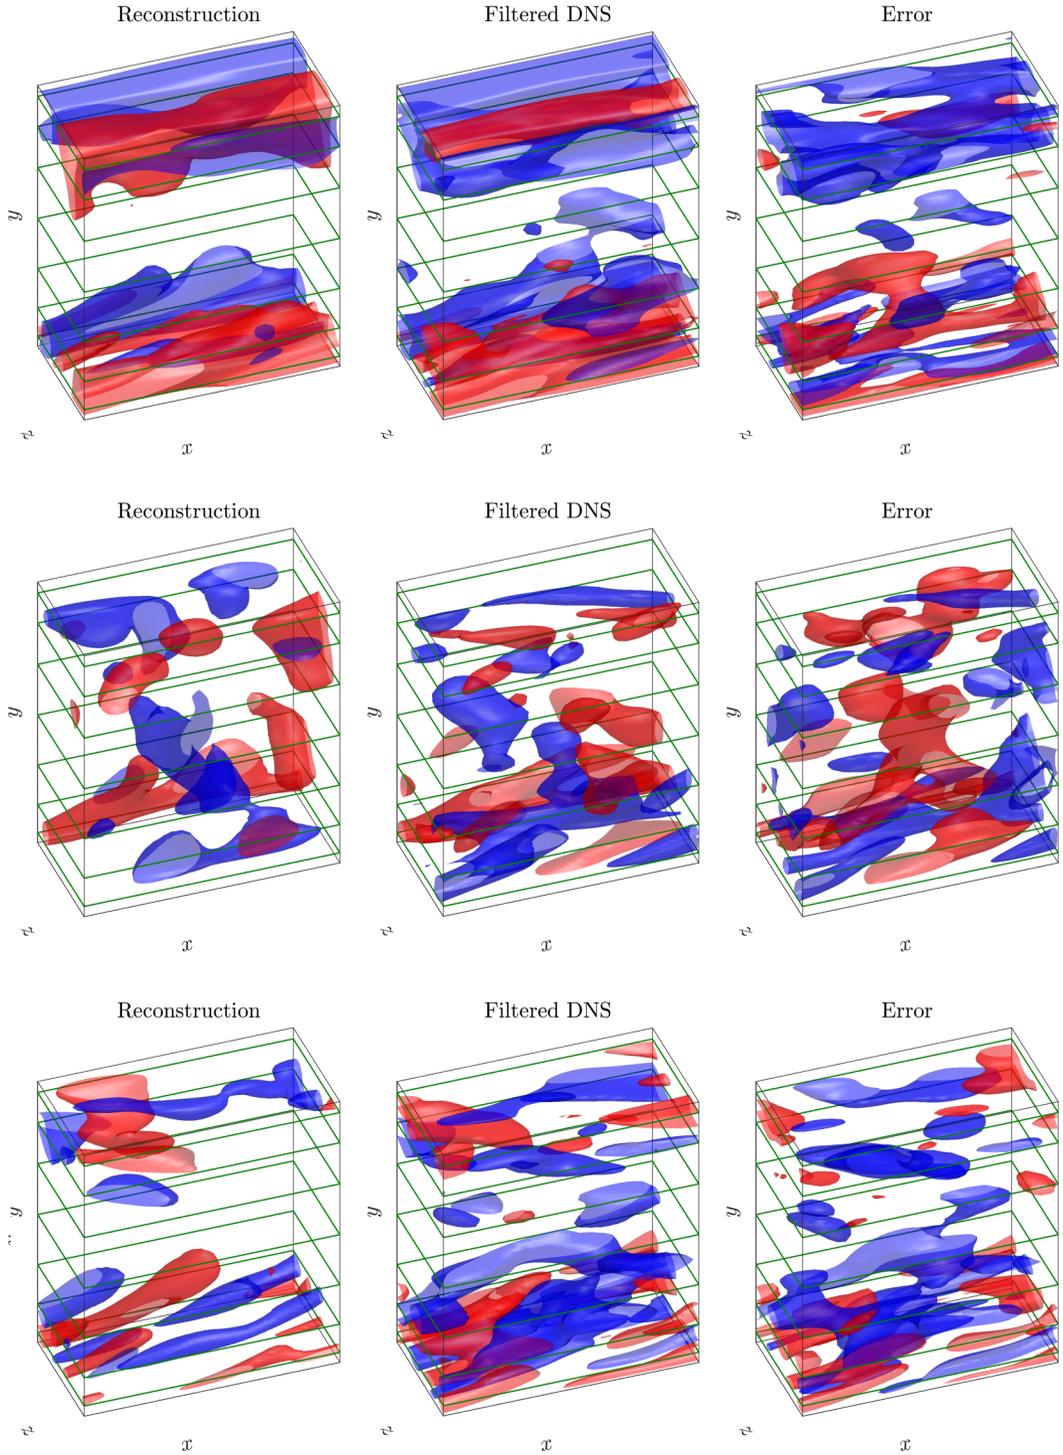


FIG. 11. Comparisons of the reconstructions, filtered DNS snapshots, and their differences in the style of Figs. 7 and 9. The comparisons here are shown for STSME with $N_{\text{plane}} = 7$ and $N_{\text{mode}} = 2$, for which $\epsilon_{\text{filt}}(t) = 0.681$ and $\epsilon_{\text{full}}(t) = 0.711$.

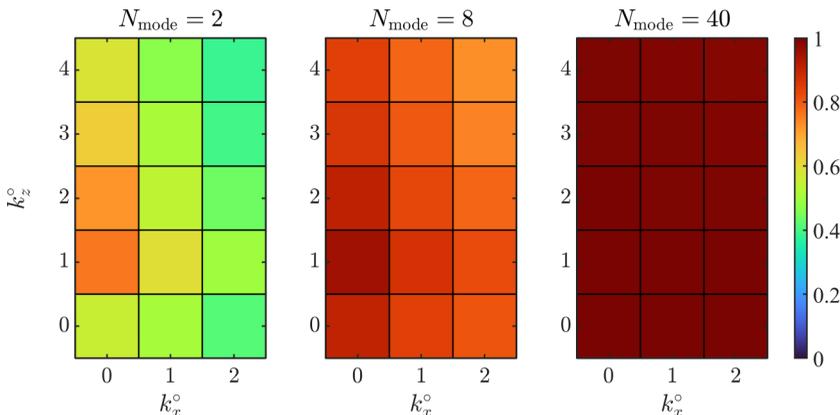


FIG. 12. The fraction of the kinetic energy captured by the first N_{mode} SPOD modes for the $N_{\text{mode}} = 2, 8,$ and 40 truncations. The fraction is shown as a function of $k_x^\circ = k_x \Delta x N_x / 2\pi \approx k_x / 3.54$ and $k_z^\circ = k_z \Delta z N_z / 2\pi \approx k_z / 7.08$, and the energies are summed over all temporal frequencies. The results are symmetric about $k_x^\circ = 0$ and $k_z^\circ = 0$.

D. Reconstruction accuracy discussion

While the STRME, SORBE, and STSME estimators considered share many themes regarding experimental design, they each have distinct use cases based on the fidelity of the available training data. When no training data are available, STRME provides an equation-based means of flow reconstruction that critically relies on the validity of the assumed mode weight statistics. The uncorrelated forcing assumption may be informative when only localized reconstructions or streamwise rms reconstructions are required. However, beyond these limited settings, second-order statistics help to inform more accurate reconstructions, especially for the wall-normal and spanwise fluctuations. When spatially limited training data are available, SORBE provides an efficient means of improving reconstruction accuracy by incorporating the data with the governing equations. Designing the estimator involves cost-accuracy tradeoffs for the fidelity of training data and the measurements. When full-field training data are available, STSME provides a data-driven framework that efficiently captures energetic flow structures on a modal basis. Practical constraints often dictate the available data and, correspondingly, the estimator used for flow reconstruction.

The linear model and training data used to form the selected estimator are critical to its behavior in constrained (e.g., experimental) settings. One common theme for the estimators we consider is that errors grow large when relatively few ($N_{\text{plane}} \leq 4$) measurements are used in conjunction with an ill-posed modal truncation (STRME and STSME) or incomplete training data (SORBE). Whereas the errors decay monotonically with N_{plane} for well-posed STSME estimators ($N_{\text{mode}} \gtrsim 40$), the STRME errors behave nonmonotonically even with all $N_{\text{mode}} = 388$ resolvent modes due to ill-posed nature of the uncorrelated forcing assumption. By contrast, the SORBE errors decay monotonically in the generalized Wiener filter limit since the statistics from training data more accurately represent the flow. These results reflect that robust, high-accuracy, full-field reconstructions require, within practical limitations, a tailored combination of (i) sufficiently high-fidelity estimators and (ii) sufficiently many measurements (e.g., $N_{\text{plane}} \geq 6$) spread throughout the channel.

1. Reconstruction performance with a single plane

While multiplane PIV is capable of simultaneously measuring multiple planes in an experimental setting, single-plane PIV represents a more conventional experimental setup. Here we discuss the errors for our $N_{\text{plane}} = 1$ generalized Wiener filter, which is particularly informative of estimator performance when measurements are isolated to one half of the channel (at $y^+ \approx 14.7$).

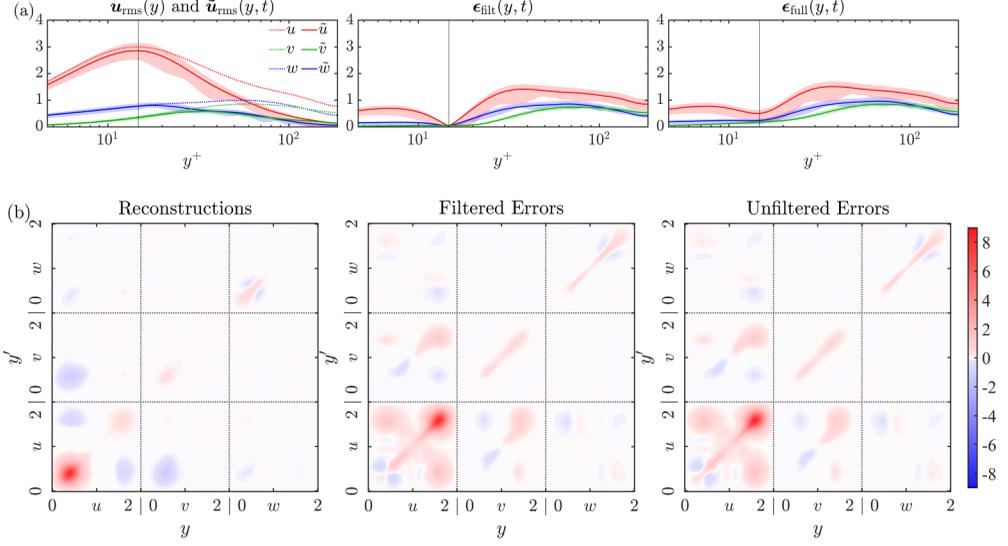


FIG. 13. Summary of the reconstruction accuracies (top: spatial profiles, bottom: cross-correlations) for the generalized Wiener filter using $N_{\text{plane}} = 1$. The plots are in the style of Figs. 6, 8, and 10.

Figure 13 shows the spatial structure of the errors associated with the $N_{\text{plane}} = 1$ generalized Wiener filter. This estimator produces highly accurate reconstructions in the vicinity of the measurement plane and, more generally, in the near-wall region. However, limited information far from the wall can be inferred from the single measurement plane. As expected by the channel symmetry, the most energetic correlations with the flow in the opposite half of the channel occur at the same wall-normal location. However, these correlations are considerably less energetic than those in the vicinity of the measurement plane. These results demonstrate that the near-wall fluctuations in each half of the channel are relatively decoupled from one another. These observations are qualitatively confirmed by the reconstruction snapshots in Fig. 14. While energetic flow structures are reconstructed for each velocity component, these structures are isolated to the bottom half of the channel. Since we place the single measurement plane at the peak of u_{rms} , it represents a logical choice for reconstructing the near-wall flow associated with the corresponding streamwise-dominated flow structures.

Beyond providing a practical benchmark for conventional experiments, the $N_{\text{plane}} = 1$ case also enables comparisons with previous studies [17,43] that employ linear estimators in similar turbulent channel flows. Consistent with these studies, we consider the error metric

$$\gamma(y, k_x, k_z) = \left(\frac{\int_{\Delta T^*} |\mathcal{F}_{x,z}\{\mathbf{u}_{\text{filt}}(\mathbf{x}, t) - \tilde{\mathbf{u}}(\mathbf{x}, t)\}|^2 dt}{\int_{\Delta T^*} |\mathcal{F}_{x,z}\{\mathbf{u}_{\text{filt}}(\mathbf{x}, t)\}|^2 dt} \right)^{1/2}, \quad (40)$$

where the integrals are taken over the testing period and $\mathcal{F}_{x,z}\{\cdot\}$ represents the spatial Fourier transform. This metric represents the distribution of relative errors over the spatial Fourier modes considered, and it is typically evaluated at estimation planes that are distinct from the measurement planes. Illingworth *et al.* [17] provides a detailed discussion of estimator performance through the lens of this metric in the context of linear models. Figure 15 shows the relative error distributions at various planes in the same half-channel as the measurement plane. These errors are qualitatively similar to those reported by Amaral *et al.* [43] using wall measurements at $\text{Re}_\tau \approx 500$ and 1000. Specifically, the present relative error levels at $y^+ \approx 28.7$ are similar to those reported previously [43] at a comparable distance (in outer units) above the measurement plane. Those error levels, which represent ORBE in a noncausal setting, were also shown [43] to be favorable to those incurred

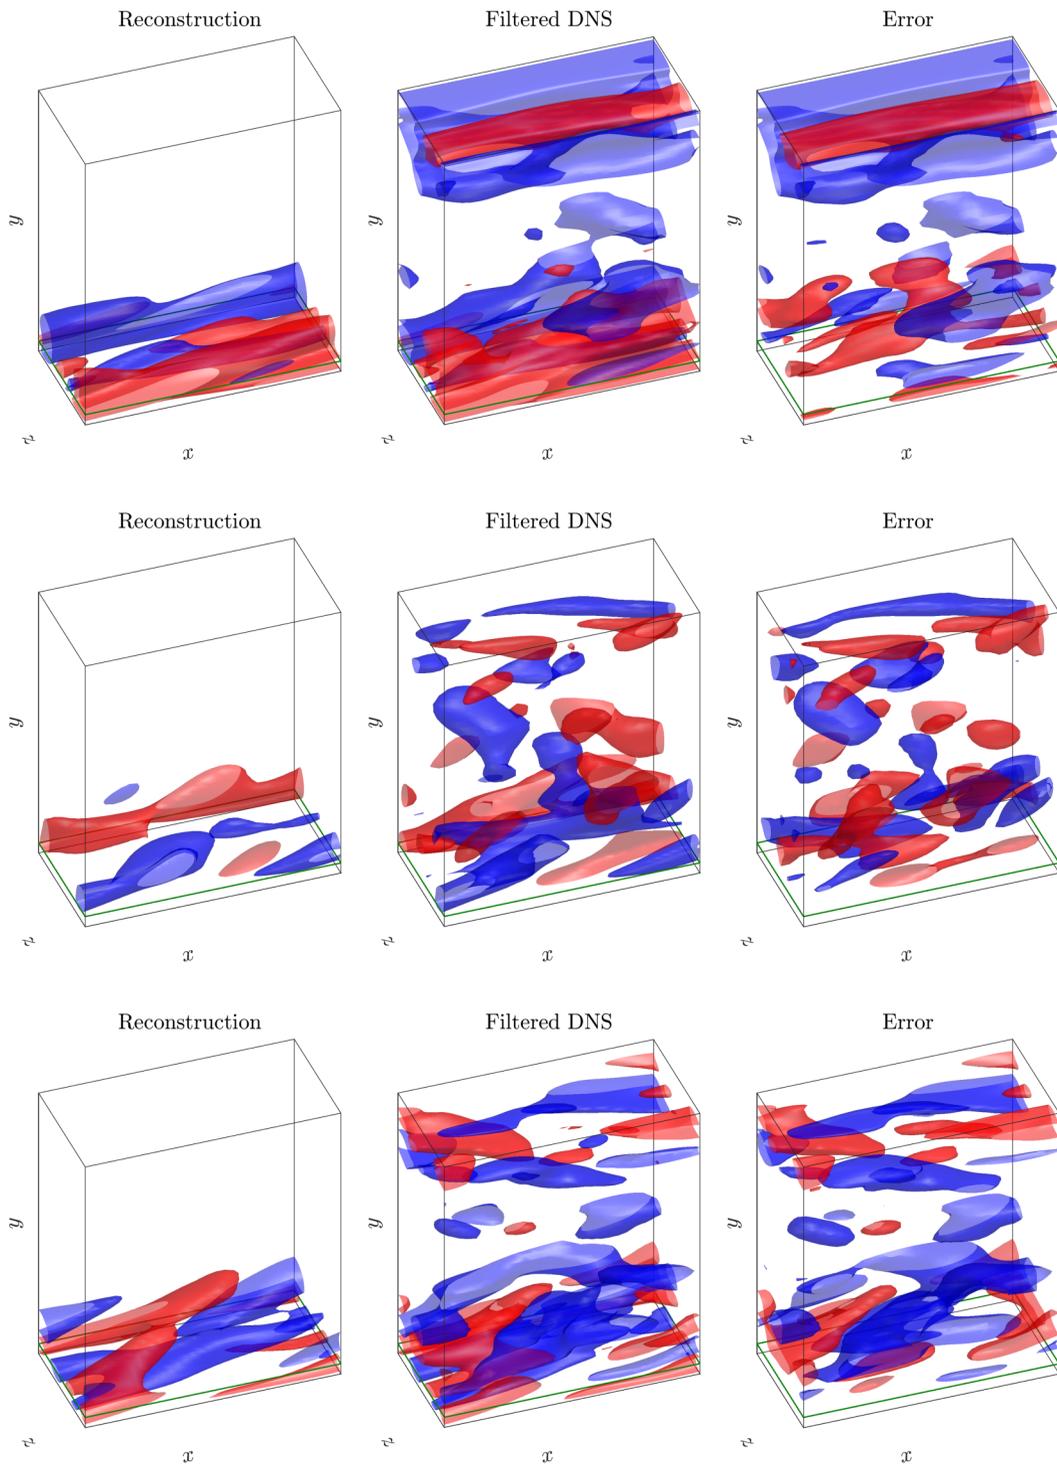


FIG. 14. Comparisons of the reconstructions, filtered DNS snapshots, and their differences in the style of Figs. 7, 9, and 11. The comparisons here are shown for the generalized Wiener filter with $N_{\text{plane}} = 1$, for which $\epsilon_{\text{filt}}(t) = 0.840$ and $\epsilon_{\text{filt}}(t) = 0.865$. Videos depicting the temporal evolutions of these comparisons are provided in the Supplemental Material [64].

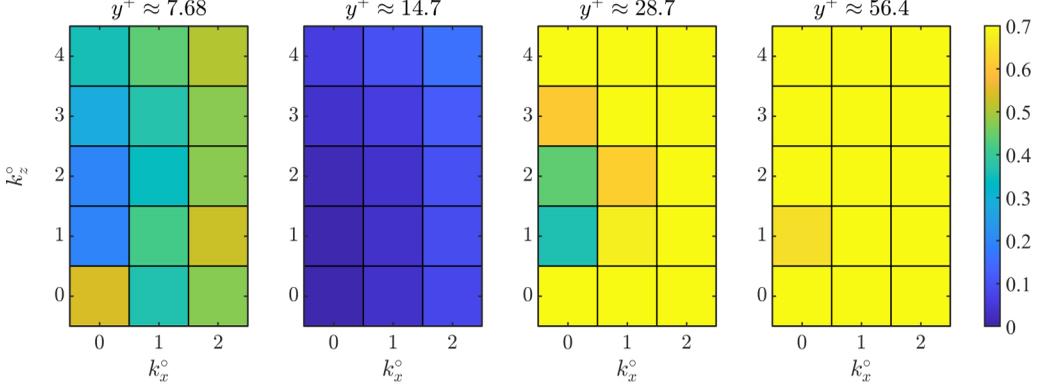


FIG. 15. Relative error levels, γ , at each estimation plane, plotted using a color axis that is consistent with previous investigations [17,43]. The normalized wave numbers, k_x^o and k_z^o , are the same as those in Fig. 12. The squared errors are averaged over positive and negative k_x^o and computed prior to averaging wall-normal velocities to the cell centers.

by the relevant estimator of Illingworth *et al.* [17]. The results of these related investigations also suggest that, while not investigated here, employing an eddy viscosity has the potential to further improve the estimation errors. However, care should be taken when comparing the present errors to those of the previous studies [17,43] since they were conducted at $\text{Re}_\tau > 186$.

2. Spatial structure of reconstruction errors

The spatial structure of the reconstruction errors, particularly in the wall-normal direction, provides insight into how streaming measurements, sensor placement, and flow statistics impact estimator performance. This structure can be formally characterized using the CSD of the estimation error, which is given by

$$\mathbb{E}[(\hat{\mathbf{u}} - \tilde{\mathbf{u}})(\hat{\mathbf{u}} - \tilde{\mathbf{u}})^H] = \mathbf{S}_{uu} - \mathbf{S}_{uy}\mathbf{T}_u^H - \mathbf{T}_u\mathbf{S}_{uy}^H + \mathbf{T}_u\mathbf{S}_{yy,n}\mathbf{T}_u^H. \quad (41)$$

As discussed previously in the context of ORBE [40,41], the error CSD expression is based on the observable subspace of linear model coefficients \mathbf{b} , as determined by the selected measurement planes. However, the expression in (41) assumes that \mathbf{S}_{uu} is fixed, consistent with the noncausal formulation of the generalized Wiener filter that formally considers a temporal Fourier transform of infinite length. Since our streaming reconstructions employ a small sliding temporal window ($\Delta T = 1$), there is a residual in the error CSD due to the temporal variations in the streaming flow statistics. During the testing period, the streaming response CSD is instantaneously given by $\mathbf{S}_{uu}^*(y, \mathbf{k}, t) = \hat{\mathbf{u}}\hat{\mathbf{u}}^H$, which is computed without averaging. This streaming CSD has a residual with respect to the true response CSD of $\boldsymbol{\varepsilon}_{uu}^*(y, \mathbf{k}, t) = \mathbf{S}_{uu}^*(y, \mathbf{k}, t) - \mathbf{S}_{uu}(y, \mathbf{k})$, and we employ analogous definitions for $\boldsymbol{\varepsilon}_{uy}^*$ and $\boldsymbol{\varepsilon}_{yy,n}^*$. The corresponding residual in the error CSD with respect to the expression in (41) is given by

$$(\hat{\mathbf{u}} - \tilde{\mathbf{u}})(\hat{\mathbf{u}} - \tilde{\mathbf{u}})^H - \mathbb{E}[(\hat{\mathbf{u}} - \tilde{\mathbf{u}})(\hat{\mathbf{u}} - \tilde{\mathbf{u}})^H] = \boldsymbol{\varepsilon}_{uu}^* - \boldsymbol{\varepsilon}_{uy}^*\mathbf{T}_u^H - \mathbf{T}_u\boldsymbol{\varepsilon}_{uy}^{*H} + \mathbf{T}_u\boldsymbol{\varepsilon}_{yy,n}^*\mathbf{T}_u^H. \quad (42)$$

Here the design of the sliding temporal window primarily impacts the contribution of $\boldsymbol{\varepsilon}_{uu}^*$. The last three terms in (42) characterize how that contribution propagates through linear estimation framework, and they depend critically on the selected measurements. Altogether, the residual in (42) characterizes how applying a formally noncausal estimator in a setting with time-varying flow statistics contributes to the estimation errors. Figure 16 depicts how the cross-correlations of the observed errors during testing are partitioned between the optimal and residual cross-correlations associated, via Parseval's theorem, with the CSDs in (41) and (42), respectively. As expected, the

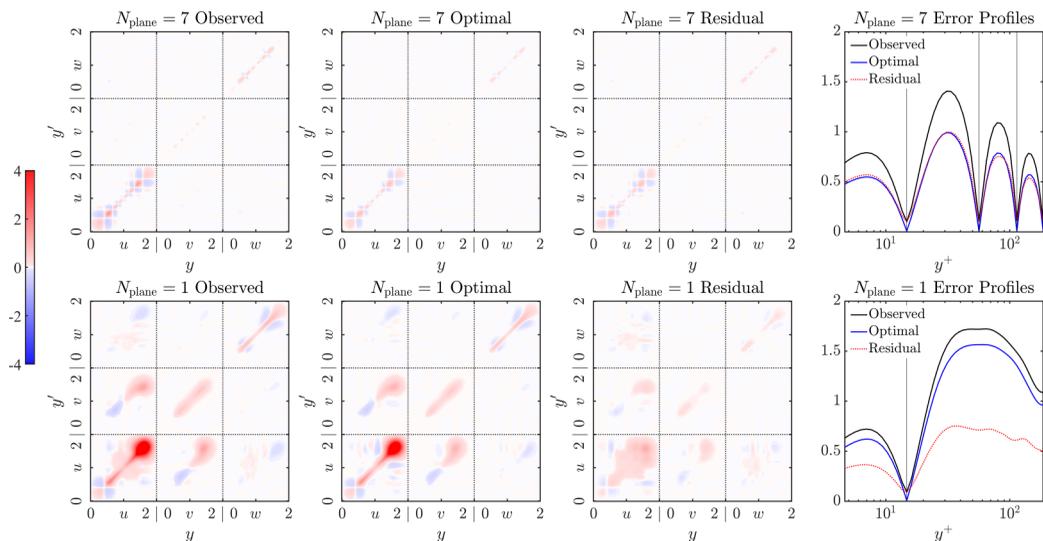


FIG. 16. Time-averaged cross-correlations of the estimation errors with respect to \mathbf{u}_{filt} during the testing period (column 1), the contributions associated with the optimal error CSDs in (41) (column 2), and those associated with the residual error CSDs in (42) (column 3). Column 4 depicts the rms error profiles, the squares of which represent the diagonals of the cross-correlations in columns 1–3, summed over all three components. These error profiles are consistent with those of ϵ_{filt} shown for generalized Wiener filter with $N_{\text{plane}} = 7$ (see Fig. 8) and $N_{\text{plane}} = 1$ (see Fig. 13).

error variance at the sensor planes directly results from the residual errors since the optimal errors decay to zero at these locations. For $N_{\text{plane}} = 7$, the variances associated with the optimal and residual errors are roughly equipartitioned away from the sensor planes. In this case, optimizing the design of the sliding Fourier transform window has the potential to further reduce errors by producing flow statistics with reduced temporal variations. By contrast, for $N_{\text{plane}} = 1$, the optimal errors dominate the residual errors far from the measurement plane. Since the true generalized Wiener filter is invariant of the linear model, this case reflects a regime in which spatially constrained measurements hinder estimator performance more than temporal uncertainty in flow statistics.

The spatial structures of the reconstruction accuracies also share various features across each linear model we consider. Far from measurements, the streamwise errors tend to grow larger than the spanwise errors, which tend to be larger than the wall-normal errors. This reflects (i) the typical magnitudes of these fluctuations and (ii) that low-error regions are more localized about measurements for u and w than for v . Physically, it reflects that the streamwise and spanwise flow structures are more localized in y , whereas the wall-normal flow structures are more persistent in y . The maximum errors (far from measurements) do not grow larger than the rms fluctuations when the estimators and measurements are of sufficiently high fidelity. Especially for SORBE and STSME, increasing the fidelity of the estimator reduces error levels throughout the channel. In the limit of full training data, the highest fidelity SORBE and STSME estimators are identical and converge to the generalized Wiener filter with the caveat of nonoptimality imposed by time-varying flow statistics.

While the full-field errors, $\epsilon_{\text{TKE}}(t)$, $\epsilon_{\text{filt}}(t)$, and $\epsilon_{\text{full}}(t)$, are informative of the global accuracy of each reconstruction method, the spatially truncated errors, $\epsilon_{\text{TKE}}^*(t)$, $\epsilon_{\text{filt}}^*(t)$, and $\epsilon_{\text{full}}^*(t)$, provide a more localized picture of reconstruction accuracy. As expected, these localized errors are considerably smaller than the full-field errors for all cases considered. However, the localized errors also tend to increase when increasing the number of sensor planes. These behaviors reflect that the larger full-field errors are concentrated far from the measurement planes and that truncated errors are computed over smaller fractions of the channel when fewer sensor planes are used. Correspondingly,

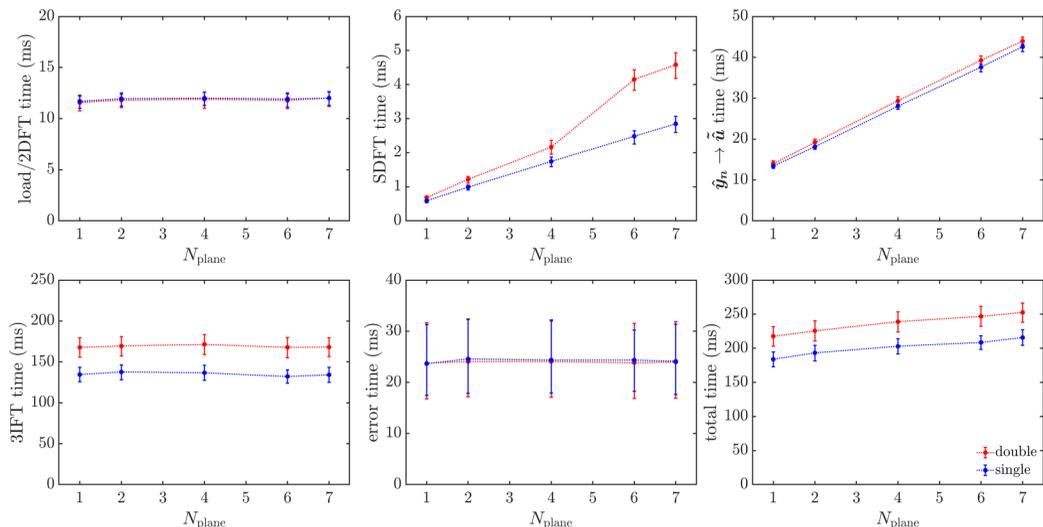


FIG. 17. Wall times required (per time step, $\Delta t = 2.86 \times 10^{-3}$) to complete each stage of the real-time flow reconstruction framework employed during the testing period. The uncertainties about the averages extend to the 20th and 80th percentiles. For $\hat{\mathbf{y}}_n \rightarrow \tilde{\mathbf{u}}$, \mathbf{T}_u is single precision for both “double” and “single” precision computations to save space. We exclude the time required to produce graphics for each snapshot.

whereas adding additional sensors can marginally increase errors near the sensor planes, this effect is outweighed globally by the reduction in errors far from the sensor planes.

In the present study, we heuristically limit the spatial integration of localized errors to within $N_{\text{local}} = 2$ cells from each measurement plane. In Appendix D we further probe how the spatial extents of the localized error domains affect reconstruction accuracy by comparing the errors, $\epsilon_{\text{filt}}^*(t)$ and $\epsilon_{\text{full}}^*(t)$, for $N_{\text{local}} = 1, 2$, and 3. The results highlight that the spatial structure of the errors can be used to limit reconstructions and thereby improve their accuracy in a manner relatively invariant to the extent of truncation.

E. Towards real-time reconstructions

In addition to reconstruction accuracy, we also evaluate the feasibility of implementing the present framework in real time (e.g., for experimental flows). We primarily focus on using the time delay required for computationally reconstructing the flow to determine relevant experimental parameters that would enable real-time reconstructions. Figure 17 provides a summary of the wall times required for each step of the reconstructions when running the (relatively unoptimized) code implementing the present frameworks on a laptop. The minimal difference between the single and double precision times when applying \mathbf{T}_u arises since \mathbf{T}_u is always single precision and only the precision of $\hat{\mathbf{y}}_n$ varies. One way that the present code is unoptimized is that converting $\tilde{\mathbf{u}}$ to $\tilde{\mathbf{u}}$ via the inverse Fourier transform (3IFT) is performed after estimating the fluctuations in Fourier space. Since these estimates are computed in parallel for each triplet, a more efficient method would simply transform and add the contributions of each triplet to the physical-space estimates as they are computed (in Fourier space). This modification will be helpful in that the 3IFT presently composes the largest fraction of the time required to reconstruct each snapshot. Further speedups may be garnered by exploiting the Hermitian symmetry of the Fourier transform when applied to real signals.

As expected, the scaling of the total reconstruction time (per snapshot) with N_{plane} results from applying the SDFT and applying \mathbf{T}_u . At the high end, the present framework reconstructs snapshots in approximately 250 ms, and we use this as our benchmark for experimental design. Figure 18

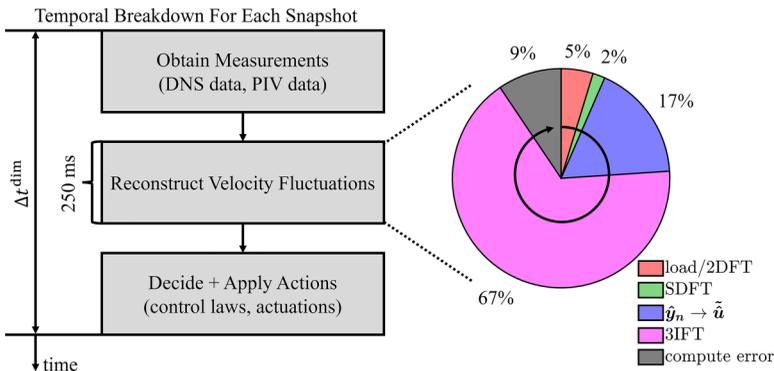


FIG. 18. Breakdown of the wall time required to achieve real-time reconstructions in an experimental setting. The pie chart shown is based on the mean wall times for double precision reconstructions using $N_{\text{plane}} = 7$, as in Fig. 17.

summarizes how this benchmark time can be used in designing an experimental protocol for real-time reconstructions. Beyond reconstruction wall time, real-time schemes must also consider the time required to (i) obtain velocity measurements (e.g., from raw images in PIV) and (ii) use the reconstructions for related tasks (e.g., flow control). Regarding (i), hardware implementations of PIV computations have computed flow fields from raw data at 15 Hz [65] and sparse processing techniques have been used to accelerate computations to up to 2000 Hz by reducing their size [66]. Nevertheless, extending real-time PIV techniques to multiplane PIV measurements may require further modifications, e.g., to accommodate multiple cameras.

We consider a hypothetical and somewhat naïve experiment by fixing $\text{Re}_\tau = 186$ and assuming flow in a water channel at room temperature (20 °C), such that $\nu = 1.00 \text{ mm}^2/\text{s}$. We further define the channel aspect ratio as $\xi = L_z/2h = 10$, where L_z is the spanwise extent of the channel, to ensure a 1D canonical mean flow. Correspondingly, the volumetric flow rate in the channel, Q_b , may be expressed in terms of the nondimensional bulk velocity, U_b , as

$$Q_b = 4\xi U_b u_\tau h^2 = 4\xi U_b \text{Re}_\tau \nu h, \quad U_b = \frac{1}{2} \int_0^2 U(y) dy. \quad (43)$$

Using the true mean profile, $U(y)$, estimated over an interval much longer than the training data, $U_b = 16.0$ for the present investigation. Similarly, the (dimensional) sampling time between snapshots required to match that of the present investigations is given by

$$\Delta t^{\text{dim}} = \frac{\Delta t}{\text{Re}_\tau} \frac{h^2}{\nu}. \quad (44)$$

We examine the quadratic, linear, and inverse dependencies of Δt^{dim} , Q_b , and u_τ , respectively, on h in Fig. 19. In order to have sufficient time to reconstruct each snapshot in real time (i.e., before the next measurement is taken), the minimum channel half-height is $h = 12.8 \text{ cm}$ and the corresponding flow rates are $O(10^4) \text{ cm}^3/\text{s}$. Both of these parameters are feasible in standard water channel configurations with sufficient pumping power, but the physical setup must be tailored to be compatible with the imaging and computation specifications. Since we fix Re_τ , the friction velocity must decrease as the size of the channel increases, and, hence, the wall roughness must be carefully designed to accommodate for the selected scale of the channel. Finally, we note that sampling every $O(100) \text{ ms}$ is a realistic benchmark for modern PIV setups. Taken together, these rough estimates indicate that the present methodology has the potential for real-time implementation in an experimental turbulent channel flow and provide motivation for evaluating the present techniques in that setting.

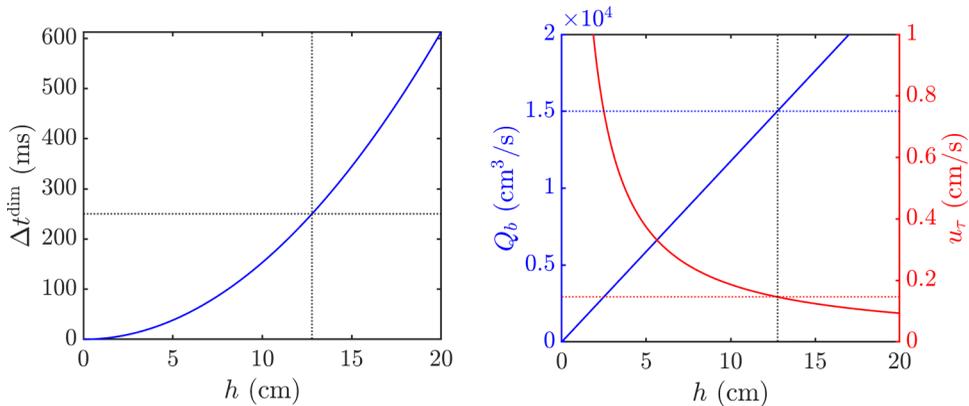


FIG. 19. Required dimensional sampling time (left), required flow rate for $\xi = 10$ (right: blue), and required friction velocity (right: red) as functions of the channel half-height. Vertical dotted lines represent the half-height corresponding to the approximate computation time to reconstruct each snapshot (250 ms).

V. CONCLUSIONS

We develop an efficient framework for streaming reconstructions of turbulent velocity fluctuations from sparse measurements with the potential for real-time implementation. In this framework, we partition time-resolved data into an initial training period and a subsequent testing period of equal length. We enable efficient, inverse-free reconstructions by first computing linear estimators from the training data and later applying them to measurements from the testing data. Various estimators are considered to address scenarios with varying degrees of available training data. Specifically, we employ streaming (i) truncated (resolvent) response mode estimation (STRME), (ii) optimal resolvent-based estimation (SORBE), and (iii) truncated SPOD mode estimation (STSME) to address the cases where the training data consists of (i) no, (ii) limited, and (iii) full-field second-order statistics.

As a prototypical example of this framework, we reconstruct snapshots from a DNS of turbulent flow in a minimal channel at $\text{Re}_\tau \approx 186$ using spatially isolated wall-parallel planes to simulate a multiplane PIV setup. We use the energetically dominant Fourier modes to efficiently represent the dynamics at each plane and evaluate their ability to reconstruct both the filtered and unfiltered snapshots. By operating in Fourier triplet space, we allow for parallel reconstructions of reduced size at each triplet. During training, we introduce blockwise inversion to efficiently and accurately compute the resolvent operator. This technique is interpretable through the governing equations in terms of input-output relationships and a modified projection. During testing, we enable efficient streaming reconstructions by recursively updating the triplet space measurements via incoming data using a temporal sliding discrete Fourier transform (SDFT).

The STRME estimators can be formed from only the governing equations and a mean profile, but rely on an unrealistic uncorrelated forcing assumption in the present formulation. They provide a simple and informative means for local reconstructions and streamwise statistical reconstructions. The SORBE estimators effectively incorporate limited training data with the governing equations to improve reconstruction accuracy. They allow for a tailored balance of the cost-accuracy tradeoffs associated with incorporating training statistics and testing measurements. The STSME estimators require full-field training statistics and do not consider the governing equations. They provide an efficient means of representing energetic flow structures on a modal basis.

With full-field flow statistics, SORBE and STSME converge to the generalized Wiener filter. This estimator produces the most accurate reconstructions in the present investigation and it has filtered and unfiltered errors of 41.5% and 52.0% (i.e., 17.2% and 27.0% of the TKE), respectively, when $N_{\text{plane}} = 7$ measurement planes are used. However, these errors are larger than those of the

formally optimal noncausal estimator since the sliding temporal window produces time-varying flow statistics. For all methods, significant reductions in error are obtained by computing the errors over a truncated wall-normal domain based on proximity to the measurements. The spatial error profiles demonstrate that this is primarily because spatially localized streamwise and spanwise flow structures are relatively well captured in these regions. The errors grow less rapidly for the wall-normal fluctuations since they are less localized in the wall-normal direction.

One promising feature of the present framework is its potential amenability to real-time reconstructions of turbulent flows in experimental settings. Even the relatively unoptimized code used in the present investigation can reconstruct each incoming snapshot in a fraction of a second (~ 250 ms). This time delay suggests that the present technique can help enable relatively accurate real-time flow estimation in laboratory settings (e.g., using a water tunnel) when paired with a multiplane PIV measurement system and a real-time PIV scheme. The time delay can be further reduced, thereby alleviating experimental design constraints, by optimizing the implementation and computational hardware used for the reconstruction scheme.

Moving forward, experimental validation is a natural next step to validate the efficacy of the present methodology in real-time settings. The velocity fluctuation statistics used for the present estimators are within reach for modern experimental measurements, and the resolvent operator relates them to the forcing statistics (as in SORBE), which are more difficult to measure directly. In estimating flow statistics, determining the required spatial resolution and optimal underlying models is important to determining the extent of training data required to produce transfer functions with sufficient fidelity. In experimental reconstructions, further considerations are required to address more complex boundary conditions and real-time processing of raw data.

Further, numerous methodological improvements are likely to yield more efficient and accurate flow reconstructions. For example, the recently developed causal ORBE framework [45] is a logical choice to employ in real-time applications. Another particularly promising avenue would be to combine frameworks to improve reconstruction accuracy with minimal training data. For example, estimators could be initialized using the STRME framework and continuously refined using flow statistics estimated from streaming measurements. Moreover, whereas the present frameworks involve computing optimal linear estimators, better performance may be obtained by introducing nonlinearity into the reconstruction formulation, either explicitly in the estimator or through, e.g., an eddy-viscosity model. Further, whereas we heuristically select measurement locations, optimizing the sensor locations to best capture the relevant dynamics in the flow is likely to improve reconstruction accuracy. While we consider a fixed measurement configuration during testing, adaptive measurement schemes may provide a means to significantly reduce measurement costs by capturing flow features of interest as they evolve. Beyond estimation, adapting the streaming framework to accommodate real-time flow control is a promising avenue of investigation. Further developments are required to address the need for in-time actuation based on real-time measurements.

ACKNOWLEDGMENTS

The support of the Office of Naval Research (ONR) under Grant No. N00014-17-1-3022 is gratefully acknowledged. R.A. was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program.

APPENDIX A: EDDY-VISCOSITY RESOLVENT OPERATOR

Following previous studies [17,43], the modified (discrete) basic linear dynamics are given by

$$L_B^t = \begin{bmatrix} L_c^t & -ik_x \partial_y v_t + \partial_y U & \mathbf{0} \\ \mathbf{0} & L_e^t - E_e & \mathbf{0} \\ \mathbf{0} & -ik_z \partial_y v_t & L_c^t \end{bmatrix}, \quad (\text{A1})$$

where $\partial_y \mathbf{v}_t$ is defined in an analogous form to $\partial_y U$. The modified diagonal subblocks are given by

$$\mathbf{L}_{c,e}^t = \mathbf{L}_{c,e} - \mathbf{E}_{c,e} - \frac{1}{\text{Re}_\tau} \frac{v_t}{\nu} \hat{\mathbf{V}}_{c,e}^2, \quad \mathbf{E}_{c,e} = \partial_y \mathbf{v}_t \partial_y, \quad (\text{A2})$$

where the discrete constituents of $\mathbf{E}_{c,e}$ are appropriately defined with respect to cell center and edge values. Here v_t is the turbulent eddy viscosity, which is given by

$$\frac{v_t}{\nu} = \frac{1}{2} \left\{ 1 + \frac{\kappa^2 \text{Re}_\tau^2}{9} (1 - \eta^2)^2 (1 + 2\eta^2)^2 \left[1 - \exp\left(\frac{\text{Re}_\tau}{A} (|\eta| - 1)\right) \right]^2 \right\}^{1/2} - \frac{1}{2}, \quad (\text{A3})$$

where $\eta = y/h$, $\kappa = 0.426$, and $A = 25.4$ [67]. Inverting the basic linear dynamics as before, now with extra terms, we express the basic resolvent as

$$\mathbf{R}_u^{B,t} = (\mathbf{L}_B^t)^{-1} = \begin{bmatrix} (\mathbf{L}_c^t)^{-1} & (\mathbf{L}_c^t)^{-1} (ik_x \partial_y \mathbf{v}_t - \partial_y U) (\mathbf{L}_e^t - \mathbf{E}_e)^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{L}_e^t - \mathbf{E}_e)^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{L}_c^t)^{-1} (ik_z \partial_y \mathbf{v}_t) (\mathbf{L}_e^t - \mathbf{E}_e)^{-1} & (\mathbf{L}_c^t)^{-1} \end{bmatrix}, \quad (\text{A4})$$

which again acts between aligned quantities in x and z for convenience. Thus, just as two inverses, \mathbf{L}_c^{-1} and \mathbf{L}_e^{-1} , are required (for staggered data) to compute the standard basic resolvent, two different inverses, $(\mathbf{L}_c^t)^{-1}$ and $(\mathbf{L}_e^t - \mathbf{E}_e)^{-1}$, are required to compute the eddy-viscosity basic resolvent. As before, using blockwise inversion to enforce continuity gives the full eddy-viscosity resolvent as

$$\mathbf{R}_u^t = \mathbf{S}^H [\mathbf{R}_u^{B,t} - \mathbf{R}_u^{B,t} \hat{\mathbf{V}} (\hat{\mathbf{V}}^T \mathbf{R}_u^{B,t} \hat{\mathbf{V}})^{-1} \hat{\mathbf{V}}^T \mathbf{R}_u^{B,t}] \mathbf{S}. \quad (\text{A5})$$

This form is analogous to the form derived for the standard resolvent operator, and therefore it may also be understood as a modified projection that acts to eliminate the pressure terms and enforce continuity.

APPENDIX B: TEMPORAL FOURIER TRANSFORM ACCURACY

We implicitly absorb into the noise vector the errors associated with computing the temporal Fourier transforms of measurements (i) using a (recursive) sliding DFT and (ii) over a finite-length window. It is therefore important to at least qualitatively characterize the validity of these methods, the former in the context of real time or streaming reconstructions and the latter in the context of transfer functions designed for noncausal estimates. Figure 20 depicts the errors associated with both of these methods for $(k_x, k_z) = (3.54, 7.08)$. The relative SDFT error plots show that the SDFT does not meaningfully contribute to the errors absorbed into the noise vector beyond those incurred by employing the `fft` over the same temporal window. The normalized temporal Fourier amplitude spectra address the validity of the finite-length temporal window. When properly normalized, the spectra computed over a temporal window of size $\Delta T = 1$ bear qualitative resemblance to those computed over $\Delta T' = 80$. Further, the variability in the amplitudes computed over subsequent smaller windows is in line with the amplitudes computed over the longer interval. This suggests that information regarding the noise incurred by truncating the integral associated with the temporal Fourier transform may be inferred to some degree. Nevertheless, to retain simplicity, we employ the simplified noise CSD model, $\hat{\mathbf{S}}_{mm} = \epsilon \mathbf{I}$, when computing the transfer functions in the present investigation.

APPENDIX C: COEFFICIENTS WITH LIMITED TRAINING DATA

The optimal choice of coefficients for flow reconstruction remains an open question. Here we empirically demonstrate that using the forcing (as in `SORBE`) produces more accurate reconstructions than using the mode weights (as in `STRME`) in the present investigation. Figure 21 compares

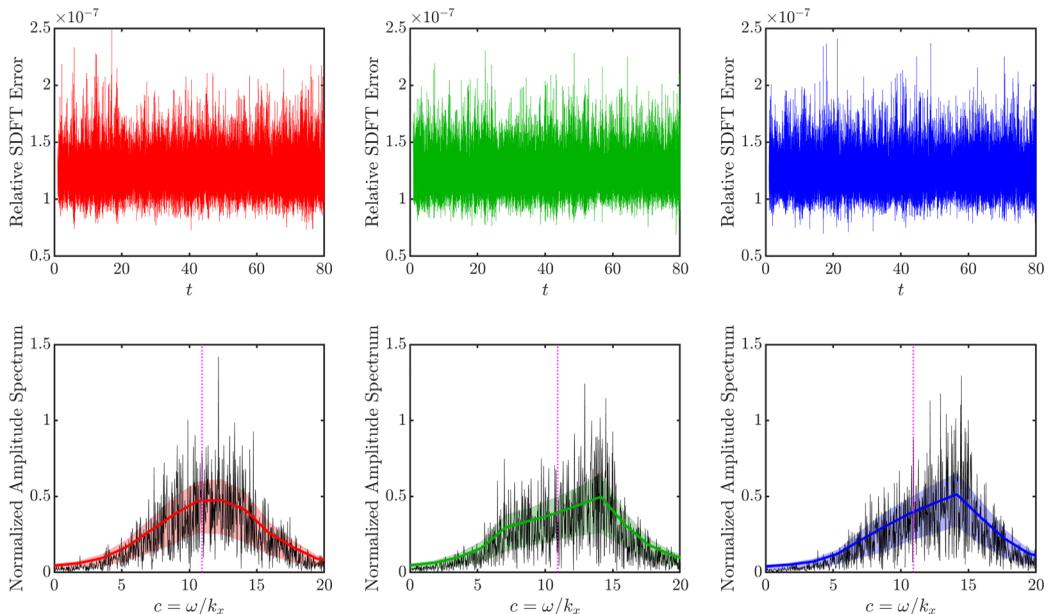


FIG. 20. Single precision relative SDFT errors (top) and normalized temporal Fourier amplitude spectra (bottom) of \hat{u} (left), \hat{v} (middle), and \hat{w} (right) computed for $(k_x, k_z) = (3.54, 7.08)$ and $y^+ \approx 14.7$ over the training period. The relative SDFT errors are evaluated with respect the `fft` coefficients computed over the same temporal window ($\Delta T = 1$). The normalized spectra show the $\Delta T' = 80$ amplitudes (black curves) and the $\Delta T = 1$ amplitudes (colored curves: square root of mean variance, shading: 20th–80th percentiles). The long-time and short-time spectra are normalized to reproduce unit signal energy, and in the latter case normalization is performed prior to computing statistics. The magenta dotted lines represent the local convection velocity, $U(y^+) \approx 10.9$.

the spatially integrated errors for estimators ($N_{\text{plane}} = 7$) that incorporate forcing and mode weight statistics computed using various (limited) amounts of training data. In both cases the coefficient statistics are estimated from the same auxiliary observations in a manner analogous to (24). The

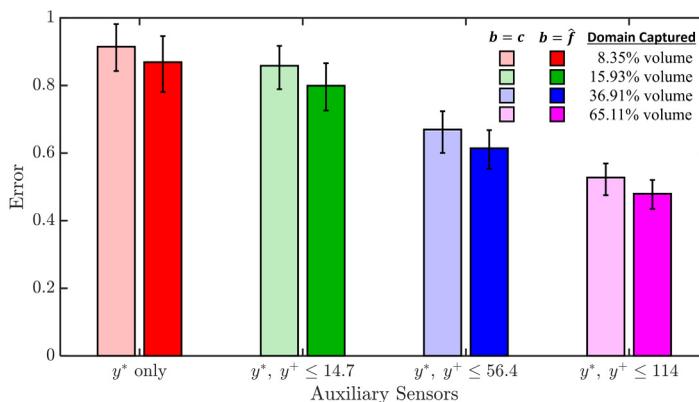


FIG. 21. Comparisons of the full-field filtered errors, $\epsilon_{\text{filt}}(t)$ (bars: square root of mean variance, error bars: 20th–80th percentiles), for reconstructions using the data-driven extension of STRME (light bars) and using SORBE (dark bars). The errors are shown for estimators ($N_{\text{plane}} = 7$) that incorporate various amounts of auxiliary training data to estimate coefficient statistics during training.

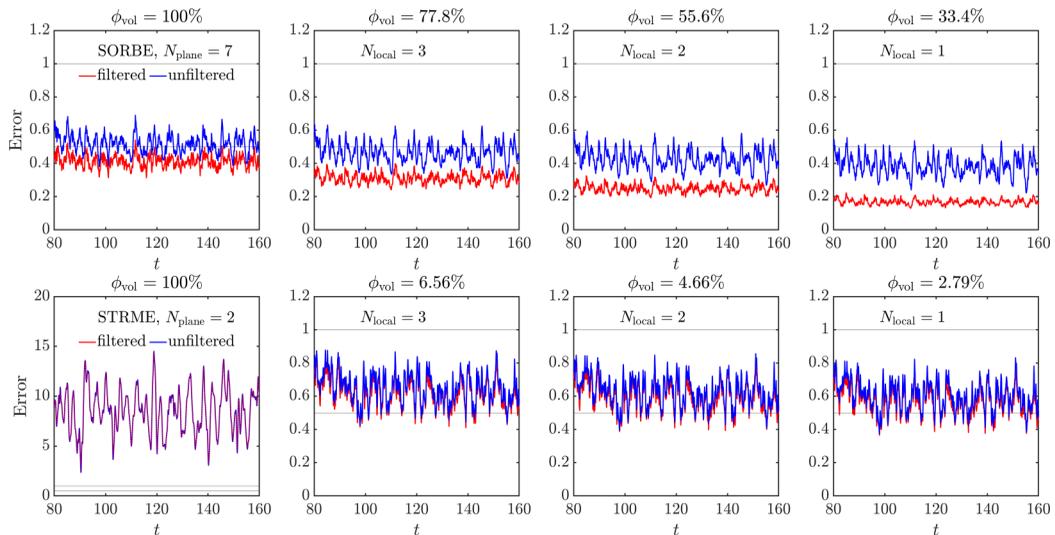


FIG. 22. Temporal evolutions of the full-field errors, $\epsilon_{\text{filt}}(t)$ and $\epsilon_{\text{full}}(t)$, and the localized errors, $\epsilon_{\text{filt}}^*(t)$ and $\epsilon_{\text{full}}^*(t)$, for $N_{\text{local}} = 3, 2$, and 1. These cases are listed such that the fraction of the channel captured by each metric, ϕ_{vol} , decreases from left to right.

results show that the forcing-based SORBE estimator produces consistently lower errors than the mode weight-based estimator representing a data-driven extension of STRME. While this validates our choice of SORBE to address the case where limited training data are available, further work is needed to find coefficients with statistics easily inferred from sparse training data.

APPENDIX D: EFFECT OF TRUNCATED ERRORS

To further probe the effect of truncating the wall-normal domain over which errors are computed, we consider truncations to within $N_{\text{local}} = 1, 2$, and 3 cells of the measurement planes. Figure 22 compares the temporal evolutions of the full-channel and localized error metrics during the testing period for selected cases. For the $N_{\text{plane}} = 7$ generalized Wiener filter, the localized reconstruction errors decrease with decreasing N_{local} , especially with respect to the filtered fluctuations. Since seven measurement planes are used, the fraction of the domain captured when integrating the errors (ϕ_{vol}) remains relatively large for all cases shown. This indicates that truncating the errors allows for improved reconstruction accuracy while retaining a large (albeit limited) domain of applicability. For the $N_{\text{plane}} = 2$ STRME estimator, the values of ϕ_{vol} are considerably smaller since fewer measurement planes are considered. These localized errors are much smaller (by an order of magnitude) than the full-field error, but apply only to considerably smaller subdomains. However, these large error reductions suggest that the $N_{\text{plane}} = 2$ STRME estimator may be suitable for localized reconstructions about each measurement plane. In this limited context, the uncorrelated forcing assumption has a relatively small impact on reconstruction accuracy.

- [1] I. Gulpe, R. Sharman, P. D. Williams, B. Zhou, G. Ellrod, P. Minnis, S. Trier, S. Griffin, S. Yum, B. Gharabaghi *et al.*, A review of high impact weather for aviation meteorology, *Pure Appl. Geophys.* **176**, 1869 (2019).

- [2] Y. Fan, G. Arwatz, T. W. Van Buren, D. E. Hoffman, and M. Hultmark, Nanoscale sensing devices for turbulence measurements, *Exp. Fluids* **56**, 138 (2015).
- [3] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns, *IEEE Control Syst. Mag.* **38**, 63 (2018).
- [4] S. L. Brunton and B. R. Noack, Closed-loop turbulence control: Progress and challenges, *Appl. Mech. Rev.* **67**, 050801 (2015).
- [5] K. Sasaki, S. Piantanida, A. V. G. Cavalieri, and P. Jordan, Real-time modelling of wavepackets in turbulent jets, *J. Fluid Mech.* **821**, 458 (2017).
- [6] I. A. Maia, P. Jordan, A. V. G. Cavalieri, E. Martini, K. Sasaki, and F. J. Silvestre, Real-time reactive control of stochastic disturbances in forced turbulent jets, *Phys. Rev. Fluids* **6**, 123901 (2021).
- [7] M. R. Abbassi, W. J. Baars, N. Hutchins, and I. Marusic, Skin-friction drag reduction in a high-Reynolds-number turbulent boundary layer via real-time control of large-scale structures, *Int. J. Heat Fluid Flow* **67**, 30 (2017).
- [8] 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).
- [9] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
- [10] 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).
- [11] O. T. Schmidt and T. Colonius, Guide to spectral proper orthogonal decomposition, *AIAA J.* **58**, 1023 (2020).
- [12] A. Ghate, A. Towne, and S. Lele, Broadband reconstruction of inhomogeneous turbulence using spectral proper orthogonal decomposition and Gabor modes, *J. Fluid Mech.* **888**, R1 (2020).
- [13] O. T. Schmidt and A. Towne, An efficient streaming algorithm for spectral proper orthogonal decomposition, *Comput. Phys. Commun.* **237**, 98 (2019).
- [14] A. Nekkanti and O. T. Schmidt, Frequency–time analysis, low-rank reconstruction and denoising of turbulent flows using SPOD, *J. Fluid Mech.* **926**, A26 (2021).
- [15] M. Chevalier, J. Høpfner, T. R. Bewley, and D. S. Henningson, State estimation in wall-bounded flow systems. Part 2. Turbulent flows, *J. Fluid Mech.* **552**, 167 (2006).
- [16] C. H. Colburn, J. B. Cessna, and T. R. Bewley, State estimation in wall-bounded flow systems. Part 3. The ensemble Kalman filter, *J. Fluid Mech.* **682**, 289 (2011).
- [17] S. J. Illingworth, J. P. Monty, and I. Marusic, Estimating large-scale structures in wall turbulence using linear models, *J. Fluid Mech.* **842**, 146 (2018).
- [18] F. Martinelli, Feedback control of turbulent wall flows, Ph.D. thesis, Politecnico di Milano, Milan, 2009.
- [19] R. J. Adrian and P. Moin, Stochastic estimation of organized turbulent structure: Homogeneous shear flow, *J. Fluid Mech.* **190**, 531 (1988).
- [20] C. Tinney, F. Coiffet, J. Delville, A. Hall, P. Jordan, and M. Glauser, On spectral linear stochastic estimation, *Exp. Fluids* **41**, 763 (2006).
- [21] B. Podvin, S. Nguiatsia, J.-M. Foucaut, C. Cuvier, and Y. Fraigneau, On combining linear stochastic estimation and proper orthogonal decomposition for flow reconstruction, *Exp. Fluids* **59**, 58 (2018).
- [22] C. He and Y. Liu, Time-resolved reconstruction of turbulent flows using linear stochastic estimation and sequential data assimilation, *Phys. Fluids* **32**, 075106 (2020).
- [23] V. Gupta, A. Madhusudanan, M. Wan, S. J. Illingworth, and M. P. Juniper, Linear-model-based estimation in wall turbulence: Improved stochastic forcing and eddy viscosity terms, *J. Fluid Mech.* **925**, A18 (2021).
- [24] A. Madhusudanan, S. J. Illingworth, and I. Marusic, Coherent large-scale structures from the linearized Navier–Stokes equations, *J. Fluid Mech.* **873**, 89 (2019).
- [25] C. Vamsi Krishna, M. Wang, M. S. Hemati, and M. Luhar, Reconstructing the time evolution of wall-bounded turbulent flows from non-time-resolved PIV measurements, *Phys. Rev. Fluids* **5**, 054604 (2020).
- [26] M. Wang, C. Vamsi Krishna, M. Luhar, and M. S. Hemati, Model-based multi-sensor fusion for reconstructing wall-bounded turbulence, *Theor. Comput. Fluid Dyn.* **35**, 683 (2021).

- [27] B. J. McKeon and A. S. Sharma, A critical-layer framework for turbulent pipe flow, *J. Fluid Mech.* **658**, 336 (2010).
- [28] B. J. McKeon, The engine behind (wall) turbulence: Perspectives on scale interactions, *J. Fluid Mech.* **817**, P1 (2017).
- [29] B. F. Farrell and P. J. Ioannou, Stochastic forcing of the linearized Navier–Stokes equations, *Phys. Fluids* **5**, 2600 (1993).
- [30] S. Symon, Reconstruction and estimation of flows using resolvent analysis and data-assimilation, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 2018.
- [31] S. Symon, D. Sipp, P. J. Schmid, and B. J. McKeon, Mean and unsteady flow reconstruction using data-assimilation and resolvent analysis, *AIAA J.* **58**, 575 (2020).
- [32] R. Moarref, A. S. Sharma, J. A. Tropp, and B. J. McKeon, Model-based scaling of the streamwise energy density in high-Reynolds-number turbulent channels, *J. Fluid Mech.* **734**, 275 (2013).
- [33] R. Moarref, M. R. Jovanović, J. A. Tropp, A. S. Sharma, and B. J. McKeon, A low-order decomposition of turbulent channel flow via resolvent analysis and convex optimization, *Phys. Fluids* **26**, 051701 (2014).
- [34] F. Gómez, H. M. Blackburn, M. Rudman, A. S. Sharma, and B. J. McKeon, A reduced-order model of three-dimensional unsteady flow in a cavity based on the resolvent operator, *J. Fluid Mech.* **798**, R2 (2016).
- [35] S. Beneddine, D. Sipp, A. Arnault, J. Dandois, and L. Lesshafft, Conditions for validity of mean flow stability analysis, *J. Fluid Mech.* **798**, 485 (2016).
- [36] L. Franceschini, D. Sipp, O. Marquet, J. Moulin, and J. Dandois, Identification and reconstruction of high-frequency fluctuations evolving on a low-frequency periodic limit cycle: Application to turbulent cylinder flow, *J. Fluid Mech.* **942**, A28 (2022).
- [37] B. Jin, S. J. Illingworth, and R. D. Sandberg, Resolvent-based approach for H_2 -optimal estimation and control: An application to the cylinder flow, *Theor. Comput. Fluid Dyn.* **36**, 491 (2022).
- [38] B. Jin, S. J. Illingworth, and R. D. Sandberg, Optimal sensor and actuator placement for feedback control of vortex shedding, *J. Fluid Mech.* **932**, A2 (2022).
- [39] In fixed-precision approximations, the rank can be tailored to efficiently satisfy a prescribed error tolerance using randomized methods [68].
- [40] A. Towne, A. Lozano-Durán, and X. Yang, Resolvent-based estimation of space–time flow statistics, *J. Fluid Mech.* **883**, A17 (2020).
- [41] E. Martini, A. V. G. Cavalieri, P. Jordan, A. Towne, and L. Lesshafft, Resolvent-based optimal estimation of transitional and turbulent flows, *J. Fluid Mech.* **900**, A2 (2020).
- [42] P. Morra, P. A. S. Nogueira, A. V. G. Cavalieri, and D. S. Henningson, The colour of forcing statistics in resolvent analyses of turbulent channel flows, *J. Fluid Mech.* **907**, A24 (2021).
- [43] F. R. Amaral, A. V. G. Cavalieri, E. Martini, P. Jordan, and A. Towne, Resolvent-based estimation of turbulent channel flow using wall measurements, *J. Fluid Mech.* **927**, A17 (2021).
- [44] P. Morra, O. Semeraro, D. S. Henningson, and C. Cossu, On the relevance of Reynolds stresses in resolvent analyses of turbulent wall-bounded flows, *J. Fluid Mech.* **867**, 969 (2019).
- [45] E. Martini, J. Jung, A. V. G. Cavalieri, P. Jordan, and A. Towne, Resolvent-based tools for optimal estimation and control via the Wiener–Hopf formalism, *J. Fluid Mech.* **937**, A19 (2022).
- [46] A. Chauhan and K. M. Singh, Recursive sliding DFT algorithms: A review, *Digital Signal Proc.* **127**, 103560 (2022).
- [47] K. Fukami, K. Fukagata, and K. Taira, Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows, *J. Fluid Mech.* **909**, A9 (2021).
- [48] However, this assumption may be violated when averaging the velocities over finite time horizons.
- [49] Another common approach to wall-normal discretization is spectral collocation.
- [50] Although not strictly written as $(i\omega\mathbf{I} - \mathbf{A})^{-1}$, we colloquially refer to \mathbf{R}_u as the resolvent operator for convenience.
- [51] K. Rosenberg and B. J. McKeon, Efficient representation of exact coherent states of the Navier–Stokes equations using resolvent analysis, *Fluid Dyn. Res.* **51**, 011401 (2019).
- [52] Although \mathbf{S}_{uu} is technically an estimate over N_{real} realizations, we assume these statistics to be converged and therefore treat such CSDs as exact.

-
- [53] Whereas the time required to generate each DNS snapshot, which is roughly $O(1)$ s here, can scale rapidly with Re_τ [69,70], this limitation is mitigated in analogous experiments.
- [54] J. Jiménez and P. Moin, The minimal flow unit in near-wall turbulence, *J. Fluid Mech.* **225**, 213 (1991).
- [55] H. J. Bae, A. Lozano-Durán, and B. J. McKeon, Nonlinear mechanism of the self-sustaining process in the buffer and logarithmic layer of wall-bounded flows, *J. Fluid Mech.* **914**, A3 (2021).
- [56] A. Liberzon, R. Gurka, and G. Hetsroni, XPIV–Multi-plane stereoscopic particle image velocimetry, *Exp. Fluids* **36**, 355 (2004).
- [57] A. Schröder and J. Kompenhans, Investigation of a turbulent spot using multi-plane stereo particle image velocimetry, *Exp. Fluids* **36**, 82 (2004).
- [58] P. A. S. Nogueira, P. Morra, E. Martini, A. V. G. Cavalieri, and D. S. Henningson, Forcing statistics in resolvent analysis: Application in minimal turbulent Couette flow, *J. Fluid Mech.* **908**, A32 (2021).
- [59] Jean Helder Marques Ribeiro, C.-A. Yeh, and K. Taira, Randomized resolvent analysis, *Phys. Rev. Fluids* **5**, 033902 (2020).
- [60] E. Martini, D. Rodríguez, A. Towne, and A. V. G. Cavalieri, Efficient computation of global resolvent modes, *J. Fluid Mech.* **919**, A3 (2021).
- [61] The pressure can alternatively be eliminated by defining a state vector of the system in terms of the wall-normal velocity and vorticity [32], but doing so complicates data processing in analogous experiments.
- [62] K. Duda, Accurate, guaranteed stable, sliding discrete Fourier transform [DSP Tips & Tricks], *IEEE Signal Proc. Mag.* **27**, 124 (2010).
- [63] Z. Kollar, F. Plesznik, and S. Trumpp, Observer-based recursive sliding discrete Fourier transform [Tips & Tricks], *IEEE Signal Proc. Mag.* **35**, 100 (2018).
- [64] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevFluids.8.064612> for videos depicting the reconstructions and their errors.
- [65] H. Yu, M. Leeser, G. Tadmor, and S. Siegel, Real-time particle image velocimetry for feedback loops using FPGA implementation, *J. Aerospace Comput. Info. Comm.* **3**, 52 (2006).
- [66] N. Kanda, C. Abe, S. Goto, K. Yamada, K. Nakai, Y. Saito, K. Asai, and T. Nonomura, Proof-of-concept study of sparse processing particle image velocimetry for real time flow observation, *Exp. Fluids* **63**, 143 (2022).
- [67] G. Pujals, M. García-Villalba, C. Cossu, and S. Depardon, A note on optimal transient growth in turbulent channel flows, *Phys. Fluids* **21**, 015109 (2009).
- [68] N. Halko, P. G. Martinsson, and J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.* **53**, 217 (2011).
- [69] D. R. Chapman, Computational aerodynamics development and outlook, *AIAA J.* **17**, 1293 (1979).
- [70] H. Choi and P. Moin, Grid-point requirements for large eddy simulation: Chapman’s estimates revisited, *Phys. Fluids* **24**, 011702 (2012).