POLITECNICO DI TORINO
Repository ISTITUZIONALE

Real-time estimation of high-resolution flow fields and reduced-order coordinates from event-based imaging velocimetry

Original
Real-time estimation of high-resolution flow fields and reduced-order coordinates from event-based imaging velocimetry / Franceschelli, L., Amico, E., Willert, C.E., Raiola, M., Cafiero, G., Discetti, S.. - In: EXPERIMENTS IN FLUIDS. - ISSN 0723-4864. - 67:12(2026). [10.1007/s00348-026-04335-7]

Availability:
This version is available at: 11583/3017218 since: 2026-10-29T19:15:30Z

Publisher:
Springer

Published
DOI:10.1007/s00348-026-04335-7

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

03 November 2026

Experiments in Fluids (2026) 67:183
https://doi.org/10.1007/s00348-026-04335-7

RESEARCH ARTICLE

Real-time estimation of high-resolution flow fields and reduced-order coordinates from event-based imaging velocimetry

Luca Franceschelli[1] · Enrico Amico[2] · Christian E. Willert[3] · Marco Raiola[1] · Gioacchino Cafiero[2] · Stefano Discetti[1]

Received: 5 May 2026 / Revised: 13 July 2026 / Accepted: 28 September 2026
© The Author(s) 2026

Abstract

A data-driven framework is proposed to estimate high-resolution (HR) velocity fields and reduced-order flow coordinates from coarse real-time event-based imaging velocimetry (rt-EBIV) measurements. Low-resolution (LR) data are retrieved through fast event analysis on a coarse grid. During offline training, the LR/HR mapping and a linear dynamical model in a proper orthogonal decomposition (POD)-based reduced-order space are identified. Online, each incoming LR snapshot is projected onto the LR basis, and the corresponding HR reduced-order coordinates are estimated and temporally regularized. Three estimators are proposed and compared: a direct Kalman filter (KF), a linear stochastic estimation (LSE) combined with Kalman filtering, and a variance-rescaled implementation (LSE+VR). The full HR velocity field is then reconstructed from the estimated reduced-order coordinates using the retained POD basis. The framework is assessed on two turbulent flows acquired by pulsed event-based imaging velocimetry (EBIV): a submerged water jet and a channel flow over a square rib. All three estimators outperform direct cubic interpolation of the LR fields, recovering physically consistent HR flow states with improved estimation of turbulent kinetic energy, spectral content, reduced-order dynamics, and temporal coherence. The linear stochastic estimator yields the lowest overall reconstruction error, while its variance-rescaled implementation provides nearly identical values and improves the recovery of fluctuation energy and higher-order content. The direct KF formulation remains highly efficient in terms of computational cost and provides the closest agreement with the HR reference in the temporal and spatial spectral analyses. The bulk of the computational cost lies in the full-field HR reconstruction, whereas the reduced-order coordinate estimation alone is negligible compared with the cost of the LR processing itself. The framework enables deliberately coarse real-time EBIV processing to be combined with reduced-order refinement, extending the effective real-time operating range towards higher update frequencies while still providing richer and dynamically more consistent HR flow representations. This makes the approach attractive for high-quality real-time flow diagnostics and as a first step towards future observer-based and closed-loop flow-control applications.

1 Introduction

Particle image velocimetry (PIV) is widely used for experimental flow diagnostics because it provides non-intrusive, spatially resolved velocity measurements over an extended field of view (Raffel et al. 2018). Compared with pointwise sensing, full-field velocimetry provides direct access to the spatial distribution of the velocity field and to quantities that depend on it, such as gradients, shear, vorticity, and multi-point correlations. This information is central to physical interpretation and particularly valuable for reduced-order modelling and observer design, since it enables state descriptions that are linked to the spatial structure of the flow rather than inferred from a small number of local probes (Glauser and George 1987; Rajaee et al. 1994; Berkooz et al. 2003). For flow-control applications, spatially distributed velocity information is particularly attractive because it supports state representations that are more directly connected to the underlying dynamics than those inferred from a limited number of local measurements (Siegel et al. 2003; Willert et al. 2010;

✉ Luca Franceschelli
luca.franceschelli@uc3m.es

[1] Department of Aerospace Engineering, Universidad Carlos III de Madrid, Avda. Universidad 30, 28911 Leganés, Madrid, Spain

[2] Department of Mechanical and Aerospace Engineering, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy

[3] DLR Institute of Propulsion Technology, German Aerospace Center, Linder Höhe, 51170 Köln, Germany

Published online: 29 October 2026

Gautier et al. 2015; Varon et al. 2019; Mccormick et al. 2024; Nonomura et al. 2025).

While PIV provides detailed spatially resolved information, it comes with limited temporal bandwidth and higher computational latency. In contrast, conventional probes (such as hot wires, pressure transducers, or microphones) are more suitable for real-time applications due to their high sampling rates, straightforward hardware integration, and minimal processing requirements. However, they provide only sparse spatial information and are often intrusive. Flow-control experiments have traditionally relied on point sensors (see e.g. Audiffred et al., 2024; Dacome et al., 2024), although their effective use generally requires careful sensor placement informed by prior knowledge of the flow, even in canonical configurations. Real-time imaging velocimetry could bridge this gap by combining the spatial information of velocity field measurements with the responsiveness needed for monitoring, estimation, and, ultimately, control.[1]

The attractiveness of implementing field information in flow-control strategies has motivated several studies aimed at advancing PIV towards real-time operation, including accelerated image-processing pipelines (Kreizer et al. 2010; Willert et al. 2010), optical-flow-based velocimetry (Gautier et al. 2015; Pimienta and Aider 2025; Bollt et al. 2025), and sparse or reduced-order reconstruction strategies (Kanda et al. 2022). However, real-time operation is generally obtained at the cost of either specialized hardware and tailored implementations, such as field-programmable gate array (FPGA)-based processing (Kreizer et al. 2010) or Graphics Processing Units (GPU) acceleration (Pimienta and Aider 2025), or a simplified measurement and processing chain. In the latter case, latency is reduced by accepting compromises such as coarser spatial resolution, fewer iterations, or lighter validation. Kanda et al. (2022) proposed sparse processing PIV (SPPIV) as a framework for real-time flow field estimation. Rather than processing the full image, cross-correlation is performed only at a small number of interrogation windows placed at optimally selected spatial locations. A pre-trained Kalman filter (KF), formulated in the proper orthogonal decomposition (POD) coefficient space, maps the resulting sparse velocity vectors onto a low-dimensional flow representation, from which the full velocity field is reconstructed. The method estimated a small number of modes owing to the limited number of simultaneous sparse measurements. The effectiveness of this approach for closed-loop flow control has been subsequently demonstrated by Nonomura et al. (2025) and Viguera et al. (2026), who employed the SPPIV-estimated velocity field as a real-time feedback signal at update frequencies of up to 2 kHz for plasma-actuator-based control. Their studies, respectively, achieved regulation of the velocity field towards prescribed targets and mitigation of flow separation over an airfoil.

Event-based vision (EBV) provides a new route to perform cost-effective flow field measurements in the kHz range. EBV replaces conventional frame-based acquisition with asynchronous brightness-change detection at the pixel level (Mahowald and Mead 1991; Lichtsteiner et al. 2008). An event sensor outputs a stream of sparse spatiotemporal events only when the local intensity change crosses a contrast threshold. Each event carries a polarity flag: positive events are triggered by a local brightness increase exceeding the contrast threshold, whereas negative events correspond to a brightness decrease. A comprehensive review of this sensing paradigm and applications is provided by Gallego et al. (2022). Compared with frame-based cameras, event sensors offer major advantages in terms of temporal responsiveness, low redundancy, reduced data transfer, and high dynamic range, making them attractive for flow diagnostics with real-time capability.

The first application of an EBV sensor to velocimetry was reported by Drazen et al. (2011), who demonstrated particle tracking in a 5-cm-diameter perspex pipe. Over the following decade, event cameras found use across a broad range of flow measurement scenarios, including Background Oriented Schlieren Method (BOS) (Shiba et al. 2024), microfluidics (Howell et al. 2020), and cardiac flows (Amico et al. 2026). Rusch (2023) and Rusch and Rösgen (2023) developed a real-time system for three-dimensional particle tracking at relatively high velocities, albeit with a hard upper bound on the number of tracks that can be reliably reconstructed. More recently, Willert (2025) proposed a dual-sensor approach that yields dense and accurate velocity fields even in high-speed flows, with demonstrated measurements in a free jet with exit velocities up to 180 m/s at laser pulse separations down to 100 ns.

Franceschelli et al. (2025) have shown that event-based imaging velocimetry (EBIV) measurements can support reduced-order flow estimation in turbulent configurations. Willert et al. (2026) demonstrated the use of EBIV for real-time jet-flow control, reaching real-time full-frame measurements at an average rate of 270 Hz with a single-pass cross-correlation analysis on pseudo-frames, thus partially sacrificing spatial resolution. This opens the following question: can coarse, data-efficient EBIV measurements be used to estimate, in real time, high-resolution velocity fields and their reduced-order coordinates, without requiring high-speed frame-based cameras, high-power laser illumination,

[1] Throughout this work, the terms *online* and *real time* are used interchangeably to denote processing performed concurrently with the experiment, at low latency. We note that, strictly speaking, real-time operation implies meeting a hard deadline dictated by the application (Stankovic 2002), whereas online operation only implies immediate, non-batch processing; the distinction is application-dependent and the relevant latency requirement here is set by the flow timescales to be resolved.

or the computational cost of a full offline-grade processing pipeline?

The present work addresses this question through a reduced-order cross-resolution estimation framework. A key point is that the low-resolution (LR) and high-resolution (HR) fields do not come from different instruments or different acquisition systems. Both datasets are extracted from the same raw EBV event stream processed according to two different strategies: a lightweight single-pass procedure representative of real-time operation provides the LR input, whereas a more accurate offline multi-frame, multi-pass procedure provides the HR reference used for training and assessment. No conventional frame-based camera data are used at any stage.

The proposed real-time estimation methodology combines this paired LR/HR processing with POD-based reduced-order modelling. In the online estimation phase, each incoming LR snapshot is projected onto the full-rank LR basis and used to estimate the corresponding reduced-order HR state. Three estimators are considered to map the LR state into the HR state: a direct Kalman Filter (KF) formulation, a Linear Stochastic Estimation (LSE)-based estimator followed by Kalman filtering, and a variance-rescaled LSE variant (LSE+VR) designed to mitigate the attenuation of poorly observable modes. The reconstructed fields are compared against the reduced-order HR reference and against direct cubic interpolation of the LR field onto the HR grid, adopted as a purely spatial baseline.

In this sense, the LR measurements can be regarded as spatially distributed, non-intrusive probes that can be combined with data-driven methods to reconstruct the corresponding HR flow field (Tu et al. 2013; Discetti et al. 2019).

The framework is evaluated on two broadband turbulent flows: a submerged water jet at $Re_D = U_j D/\nu \approx 6,000$, and a channel flow over a square rib at $Re_H = U_{ch} H/\nu \approx 1,500$. Here, U_j is the jet bulk velocity, U_{ch} is the asymptotic channel-flow velocity, D is the nozzle exit diameter, H is the rib height, and ν is the kinematic viscosity. These two cases were selected because they exhibit energetically distributed, multi-scale dynamics and therefore provide a demanding testbed for cross-resolution estimation. In addition to instantaneous fields, the analysis includes mean and second-order statistics, temporal and spatial spectra, reduced-order modal dynamics, temporal autocorrelation, overall reconstruction error, and computational cost.

The complete measurement and estimation workflow presented in this work is released as *VIBEstream*, an open-source Python package providing a simple interface for event-camera recording, live streaming, quick PIV-style processing, and real-time EBIV with direct velocity field output (see the code availability statement).

The remainder of the paper is organized as follows. Section 2 introduces the common reduced-order formulation and the three estimation methods. Section 3 describes the experimental datasets and the EBIV processing chain used to generate the paired LR and HR fields. Section 4 presents the reconstruction results and the computational cost analysis. Finally, Sect. 5 summarizes the work and discusses the practical implications, limitations, and possible extensions of the proposed framework.

2 Methodology

This section describes the three estimation methodologies considered in the present work. Although they differ in the way the HR coefficients are estimated and temporally regularized, all three methods share the same offline/online framework and use the temporal coefficients from LR measurements as input.

For clarity, the main terminology used throughout the paper is defined here. The *latent coordinates* are the POD temporal coefficients, i.e. the modal amplitudes. The *reduced-order coordinates* denote the subset of these coefficients retained after truncation of the POD basis. Accordingly, the *latent space*, or equivalently the *reduced-order space*, is the subspace spanned by the retained POD modes. In this work, *state estimation* refers to the estimation of the retained POD coefficients from coarse real-time measurements. Finally, the *modal energy* denotes the POD modal energy, defined as the squared singular value associated with each mode.

The workflow pipeline is reported in Fig. 1. The main structure is divided into two phases:

- **Offline training**: event sequences are analysed both with a simple coarse analysis based on single-pass cross-correlation with non-overlapping sampling window (LR dataset) and with an advanced multi-pass multi-frame process (HR dataset). A mapping is then established between the LR and HR data with the aid of reduced-order coordinates.
- **Online estimation** : the stream of events is processed with the same quick evaluation process used to obtain the LR dataset. The data are then used to estimate the HR velocity fields using the mapping learned during training.

More explicitly, in the offline phase (top, grey dashed frame of Fig. 1) the recorded event data are processed by both pipelines: the single-pass LR processing and the multi-pass, multi-frame HR processing, the latter followed by the rank truncation discussed in Sect. 2.1. POD applied to the two datasets yields the LR basis and singular values $\mathbf{\Phi}^{\mathrm{LR}}\mathbf{\Sigma}^{\mathrm{LR}}$ with the associated coefficients $\mathbf{\Psi}^{\mathrm{LR}}$ and their truncated HR counterparts. From the paired coefficient histories, the LR/HR mapping operator and the dynamics model matrices are identified. In the online phase (bottom, blue dashed

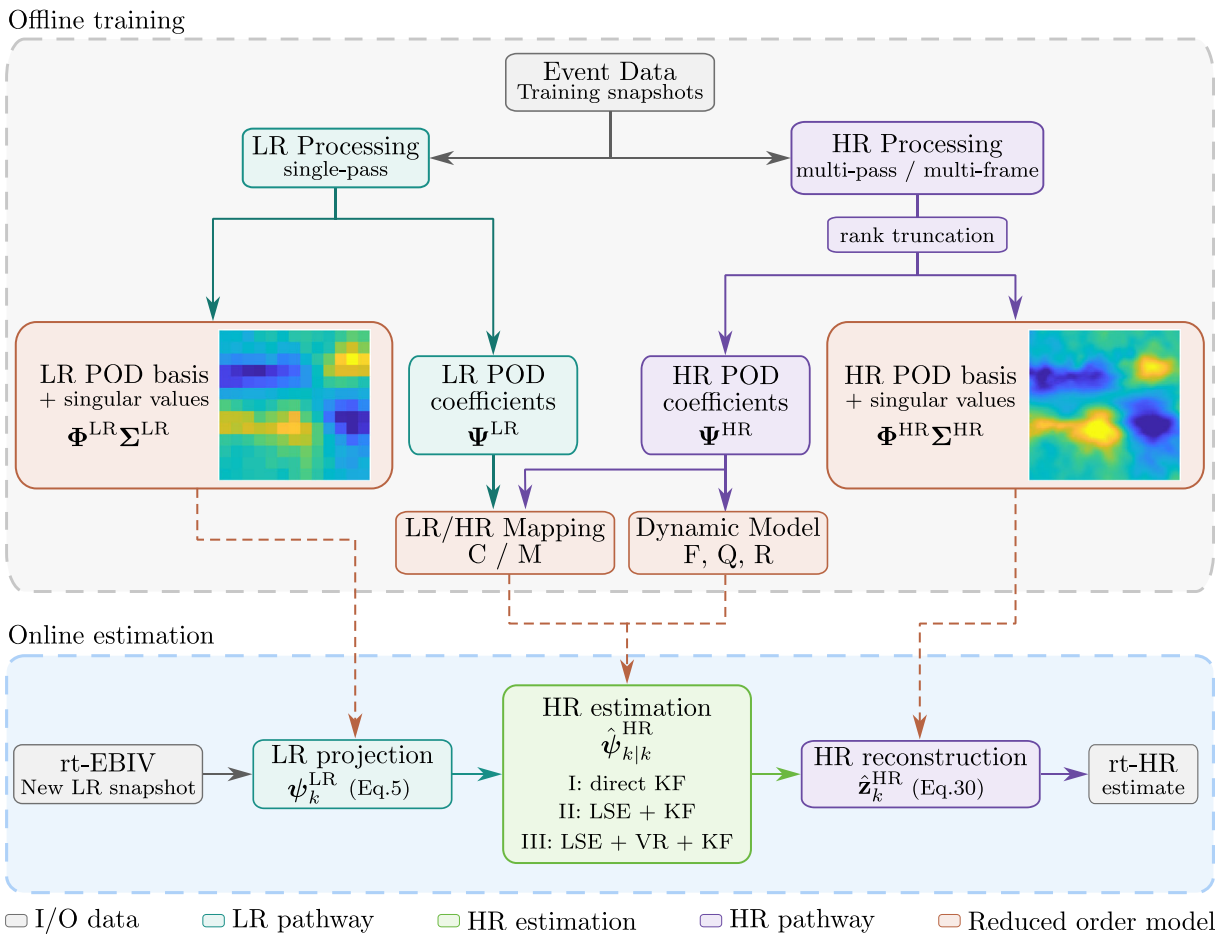


Fig. 1 Offline training (top, grey dashed frame) and online estimation (bottom, blue dashed frame) pipeline. The colour coding distinguishes input/output data (grey), the LR pathway (teal), the HR pathway (purple), the reduced-order model operators (orange), and the HR estimation stage (green); dashed orange arrows denote the operators identified offline and deployed online. The abbreviation "rt" stands for real-time

frame), each new LR snapshot delivered by the real-time EBIV processing is projected onto the LR basis, Eq. (5); the reduced-order HR state $\hat{\boldsymbol{\psi}}_{k|k}^{\mathrm{HR}}$ is then estimated through one of the three methods of the following sections, using the operators learned offline (dashed orange arrows in the figure); finally, when the full field is required, the HR snapshot is reconstructed through the retained HR basis.

From a practical standpoint, once the training is completed, the online stage operates as a self-contained estimator: each incoming LR snapshot is converted, within the available time budget, into an estimate of the reduced-order HR coordinates and, if required, into the corresponding full HR field. In a prospective observer-based or closed-loop flow-control architecture, the estimated reduced-order coordinates would play the role of the plant state: they could be fed directly to a state-feedback controller or serve as the measurement input of a higher-level observer, without requiring the full-field reconstruction in the loop.

A key point of the present framework is that the LR and HR velocity fields do not originate from different measurement systems. In both datasets, the same raw EBV event stream is processed through two distinct velocimetry pipelines of different computational cost and accuracy. The first is a lightweight single-pass procedure designed to emulate a realistic real-time implementation and provides the LR input to the estimator. The second is a more accurate offline multi-frame, multi-pass procedure and provides the HR reference fields used for training, validation, and performance assessment. No conventional frame-based camera data are used at any stage of the present work.

The main hypothesis is that a LR-HR mapping can be established by using a compact mapping in latent coordinates. This approach is in line with previous works exploiting reduced-order representations and/or nonlinear mappings for super-resolution (Deng et al. 2019; Fukami et al. 2021; Cortina-Fernández et al. 2021; Tirelli et al. 2023).

In this section, the common formulation is introduced first, followed by the specific description of each method. First, the POD implementation is briefly summarized, as it is used to identify the temporal coordinates for the mapping between LR and HR measurements, and for the state transitions. Later, the implementation of state observers for HR full-state estimation is presented.

2.1 Proper orthogonal decomposition and shared framework

Following the snapshot method introduced by Sirovich (1987), POD is applied to the HR training dataset in fluctuation form. We denote the instantaneous velocity field as

$$\mathbf{U}(\mathbf{x},t) = \overline{\mathbf{U}}(\mathbf{x}) + \mathbf{u}(\mathbf{x},t)\,, \tag{1}$$

where $\overline{\mathbf{U}}(\mathbf{x}) = \langle \mathbf{U}(\mathbf{x},t)\rangle_T$ denotes the temporal mean field and $\mathbf{u}(\mathbf{x},t)$ the corresponding fluctuation field according to the Reynolds decomposition. In PIV measurements, each snapshot samples the c components of the fluctuating velocity field over a grid of sampling points covering the measurement domain. Let N_p^{HR} and N_p^{LR} denote the number of spatial sampling points of the HR and LR velocity fields, respectively. Then, $\mathbf{z}_k^{\mathrm{HR}} \in \mathbb{R}^{c\,N_p^{\mathrm{HR}}}$ and $\mathbf{z}_k^{\mathrm{LR}} \in \mathbb{R}^{c\,N_p^{\mathrm{LR}}}$ denote the corresponding k-th HR and LR fluctuation snapshots. Each snapshot is formed by stacking all the velocity components over all the available sampling points into a single column vector. Denoting by $\mathbf{Z}^{\mathrm{HR}} \in \mathbb{R}^{cN_p^{\mathrm{HR}}\times N_t}$ the matrix collecting the N_t HR training snapshots, its economy-sized singular value decomposition is

$$\mathbf{Z}^{\mathrm{HR}} = \left[\mathbf{z}_1^{\mathrm{HR}},\ldots,\mathbf{z}_{N_t}^{\mathrm{HR}}\right] = \Phi^{\mathrm{HR}}\Sigma^{\mathrm{HR}}\Psi^{\mathrm{HR}^{\mathsf{T}}}, \tag{2}$$

where $\Phi^{\mathrm{HR}} \in \mathbb{R}^{cN_p^{\mathrm{HR}}\times m}$ contains the spatial modes, $\Psi^{\mathrm{HR}} \in \mathbb{R}^{N_t\times m}$ the temporal modes, $\Sigma^{\mathrm{HR}} \in \mathbb{R}^{m\times m}$ the singular values, and $m \le \min(N_t, cN_p^{\mathrm{HR}})$ the maximum rank of $\mathbf{Z}^{\mathrm{HR}}$. A reduced-order representation of the HR field can be obtained through a Low-Order Reconstruction (LOR) by truncating the HR latent space from rank m to rank r, retaining only the first r modes:

$$\mathbf{Z}^{\mathrm{HR}} \approx \boldsymbol{\Phi}_r^{\mathrm{HR}}\boldsymbol{\Sigma}_r^{\mathrm{HR}}\boldsymbol{\Psi}_r^{\mathrm{HR}^{\mathsf{T}}}, \tag{3}$$

where $\boldsymbol{\Phi}_r^{\mathrm{HR}} \in \mathbb{R}^{cN_p^{\mathrm{HR}}\times r}$ contains the retained spatial modes, $\boldsymbol{\Sigma}_r^{\mathrm{HR}} \in \mathbb{R}^{r\times r}$ is the corresponding diagonal matrix of singular values, and $\boldsymbol{\Psi}_r^{\mathrm{HR}} \in \mathbb{R}^{N_t\times r}$ contains the associated temporal modes. Hereafter, $\boldsymbol{\psi}_k^{\mathrm{HR}} \in \mathbb{R}^r$ denotes the reduced HR coefficient vector associated with the k-th snapshot, i.e. the k-th row of $\boldsymbol{\Psi}_r^{\mathrm{HR}}$ transposed to a column vector.

The choice of the truncation rank involves a delicate trade-off between preserving physically meaningful flow structures and rejecting noise, which typically contaminates the higher-order modes. In this sense, previous studies have proposed criteria based on the modal amplitude relative to the measurement noise (Epps and Techet 2010), on the degree of spatial organization retained by the modes (Brindise and Vlachos 2017), and on the progressive loss of modal reliability as noise corruption increases (Epps and Krivitzky 2019). In this work, the truncation rank r is selected using the *elbow criterion*, restated by Raiola et al. (2015) as the relative decrease rate of the reconstruction error at the i-th POD mode, $F(i)$, rather than on a prescribed retained energy threshold. The threshold $F(i) = 0.999$ is adopted to determine the number of retained HR modes. This threshold identifies the mode index beyond which adding one more mode reduces the reconstruction error by less than 0.1% in relative terms, and the additional modes are expected to be dominated by measurement noise. Other values were also considered: stricter truncations discard part of the energetic broadband content of the present flows, whereas milder ones mostly add noise-dominated modes with no appreciable benefit on the estimation. The relatively high retained ranks ($r \approx 180$–190) reflect the broadband, multi-scale character of the two datasets, whose slowly decaying modal energy distribution (see Sect. 4) requires many modes to represent the resolved dynamics, rather than the specific threshold choice. Note also that the POD truncation acts at the same time as an effective denoising stage, since measurement noise is mostly scattered across the discarded low-energy modes (Raiola et al. 2015; Brindise and Vlachos 2017).

An analogous POD analysis is applied to the LR training dataset. Denoting by $\mathbf{Z}^{\mathrm{LR}} \in \mathbb{R}^{cN_p^{\mathrm{LR}}\times N_t}$ the matrix of rank n collecting the LR training snapshots, its singular value decomposition yields

$$\mathbf{Z}^{\mathrm{LR}} = \boldsymbol{\Phi}^{\mathrm{LR}}\boldsymbol{\Sigma}^{\mathrm{LR}}\boldsymbol{\Psi}^{\mathrm{LR}^{\mathsf{T}}}. \tag{4}$$

In contrast to the HR side, no truncation is applied to the LR latent space, and the full numerical rank n is retained.

During the validation and online reconstruction stage, i.e. for the snapshots belonging to the testing subset, each incoming LR snapshot $\mathbf{z}_k^{\mathrm{LR}}$ is projected onto the previously identified full-rank LR basis to obtain the coefficient vector

$$\boldsymbol{\psi}_k^{\mathrm{LR}} = \boldsymbol{\Sigma}^{\mathrm{LR}^{-1}}\boldsymbol{\Phi}^{\mathrm{LR}^{\mathsf{T}}}\mathbf{z}_k^{\mathrm{LR}} \in \mathbb{R}^n, \tag{5}$$

which constitutes the common input to all three estimation methodologies described below.

In the following, we consider two-component velocity measurements, i.e. $c = 2$, although the extension to volumetric three-component data is straightforward.

2.2 Estimation of HR states from LR measurements

In this section, three estimators will be described:

- **Direct Kalman filter estimation**: the HR temporal coefficients are obtained using directly the LR data as observables for the Kalman filter.
- **LSE-based estimation with Kalman filtering**: the HR coefficients are first estimated with an LSE from the LR coefficients, and then a Kalman filter is applied as a regularizer.
- **Variance-rescaled LSE-based estimation**: based on the same strategy of the previous method, corrected with a mode-variance correction before the Kalman filtering to compensate for systematic attenuation of the least observable HR modes.

The three estimators differ in the way the reduced HR coefficients are inferred from the LR measurements. However, all of them rely on the same linear reduced-order state-transition model of the HR coefficients. This common dynamical model is introduced first, and the three estimation strategies are described subsequently.

For the estimator identification, it is convenient to collect the reduced coordinates at all time instants as column-wise data matrices,

$$\begin{aligned} \mathbf{A}^{\mathrm{HR}} &\equiv \left[\boldsymbol{\psi}_1^{\mathrm{HR}}, \ldots, \boldsymbol{\psi}_{N_t}^{\mathrm{HR}}\right] = \boldsymbol{\Psi}_r^{\mathrm{HR}\mathsf{T}} \in \mathbb{R}^{r\times N_t}, \\ \mathbf{A}^{\mathrm{LR}} &\equiv \left[\boldsymbol{\psi}_1^{\mathrm{LR}}, \ldots, \boldsymbol{\psi}_{N_t}^{\mathrm{LR}}\right] = \boldsymbol{\Psi}^{\mathrm{LR}\mathsf{T}} \in \mathbb{R}^{n\times N_t}. \end{aligned} \tag{6}$$

These matrices are introduced only as a compact notation for the state estimation stage: each column corresponds to one time sample of the reduced HR state or LR measurement.

2.2.1 Common state-transition model

During online estimation, a state-transition model is used to regularize the output. Considering the short prediction horizon (only one step ahead in time-resolved measurements), a linear discrete one-step model is adopted for the reduced HR state. Specifically, the retained HR temporal coefficients are assumed to evolve according to a linear state-transition matrix $\mathbf{F} \in \mathbb{R}^{r\times r}$, such that

$$\boldsymbol{\psi}_{k+1}^{\mathrm{HR}} \approx \mathbf{F}\,\boldsymbol{\psi}_k^{\mathrm{HR}}. \tag{7}$$

To estimate $\mathbf{F}$, two matrices collecting consecutive reduced-order snapshots are defined as

$$\begin{aligned} \mathbf{A}^{\mathrm{HR},-} &= \left[\boldsymbol{\psi}_1^{\mathrm{HR}} \;\cdots\; \boldsymbol{\psi}_{N_t-1}^{\mathrm{HR}}\right] \in \mathbb{R}^{r\times(N_t-1)}, \\ \mathbf{A}^{\mathrm{HR},+} &= \left[\boldsymbol{\psi}_2^{\mathrm{HR}} \;\cdots\; \boldsymbol{\psi}_{N_t}^{\mathrm{HR}}\right] \in \mathbb{R}^{r\times(N_t-1)}, \end{aligned} \tag{8}$$

The matrix $\mathbf{F}$ is then obtained in the least-squares sense as

$$\mathbf{F} = \mathbf{A}^{\mathrm{HR},+}\left(\mathbf{A}^{\mathrm{HR},-}\right)^{\dagger}, \tag{9}$$

where $(\cdot)^{\dagger}$ denotes the Moore–Penrose pseudoinverse. Therefore, $\mathbf{F}$ represents the best-fit linear evolution operator of the reduced HR state over one time step.

2.2.2 Method I—Direct Kalman filter estimation

The first estimator directly infers the reduced HR state from the LR coefficient vector through a linear state-space model formulated entirely in coefficient space. Since the estimation is directly embedded in the Kalman filtering process, this method will be referred to as KF.

The reduced HR state is assumed to evolve according to

$$\boldsymbol{\psi}_k^{\mathrm{HR}} = \mathbf{F}\,\boldsymbol{\psi}_{k-1}^{\mathrm{HR}} + \mathbf{w}_k, \tag{10}$$

where $\mathbf{w}_k$ denotes the process noise of the Kalman filter state model, accounting for unmodelled dynamics and mismatch between the true reduced-order evolution and its one-step linear approximation through $\mathbf{F}$. It is assumed to be zero mean with covariance $\mathbf{Q}$, estimated offline from the residuals of the identified linear model. The measurement model relates the LR coefficients vector to the reduced HR state as

$$\boldsymbol{\psi}_k^{\mathrm{LR}} = \mathbf{C}\,\boldsymbol{\psi}_k^{\mathrm{HR}} + \boldsymbol{\eta}_k, \tag{11}$$

where $\mathbf{C} \in \mathbb{R}^{n\times r}$ is a linear measurement operator mapping HR coefficients to LR coefficients, and $\boldsymbol{\eta}_k$ denotes the measurement noise, assumed zero mean with covariance $\mathbf{R}$.

The comparison between the HR prediction and the LR measurement is deliberately performed in the LR coefficient space. Indeed, mapping a LR measurement back to the HR space before comparison would amount to solving an inverse problem, which is intrinsically ill-posed because the LR data do not contain the missing high-frequency content. By contrast, projecting the predicted HR state into the LR space through $\mathbf{C}$ is a well-posed forward operation. The resulting mismatch is then mapped back to the HR state through the covariance-weighted Kalman gain, ensuring a stable and consistent correction step.

The operator $\mathbf{C}$ is identified offline from the training data by ridge-regularized least squares,

$$\mathbf{C} = \mathbf{A}_{\mathrm{train}}^{\mathrm{LR}}\mathbf{A}_{\mathrm{train}}^{\mathrm{HR}\mathsf{T}}\left(\mathbf{A}_{\mathrm{train}}^{\mathrm{HR}}\mathbf{A}_{\mathrm{train}}^{\mathrm{HR}\mathsf{T}} + \lambda\mathbf{I}_r\right)^{-1}, \tag{12}$$

where λ is a small regularization parameter, here set to $\lambda = 10^{-12}$.

The process noise covariance $\mathbf{Q}$ is estimated empirically from the residuals of the one-step linear model on the training

set. Specifically, the matrix

$$\mathbf{E}_{\text{model}} = \mathbf{A}^{\text{HR},+} - \mathbf{F}\,\mathbf{A}^{\text{HR},-} \tag{13}$$

collects, column by column, the sample residuals of the reduced-order state equation, i.e. the observed realizations of the process noise term $\mathbf{w}_k$ over the training set. The corresponding covariance estimate is then

$$\mathbf{Q} = \frac{1}{N_t - 1}\mathbf{E}_{\text{model}}\mathbf{E}_{\text{model}}^{\mathsf{T}}\,. \tag{14}$$

Similarly, the measurement noise covariance $\mathbf{R}$ is estimated from the residuals of Eq. (11) over the validation subset. Denoting by $\mathbf{A}_{\text{val}}^{\text{LR}}$ and $\mathbf{A}_{\text{val}}^{\text{HR}}$ the matrices collecting the corresponding LR and HR coefficients vectors on that subset, one defines

$$\mathbf{E}_{\text{meas}} = \mathbf{A}_{\text{val}}^{\text{LR}} - \mathbf{C}\,\mathbf{A}_{\text{val}}^{\text{HR}}\,. \tag{15}$$

The matrix $\mathbf{E}_{\text{meas}}$ therefore collects, column by column, the sample residuals of the measurement equation and is used as an empirical collection of realizations of the measurement noise term $\boldsymbol{\eta}_k$. The corresponding covariance estimate is

$$\mathbf{R} = \frac{1}{N_{\text{val}} - 1}\mathbf{E}_{\text{meas}}\mathbf{E}_{\text{meas}}^{\mathsf{T}}\,. \tag{16}$$

Once $\mathbf{F}$, $\mathbf{C}$, $\mathbf{Q}$, and $\mathbf{R}$ have been identified, the reduced HR state is estimated online through the classical Kalman prediction-update procedure. At each k-th timestep, the Kalman gain is updated as

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\,\mathbf{C}^{\mathsf{T}}\left(\mathbf{C}\,\mathbf{P}_{k|k-1}\,\mathbf{C}^{\mathsf{T}} + \mathbf{R}\right)^{-1}\,, \tag{17}$$

where $\mathbf{P}_{k|k-1}$ denotes the covariance matrix of the *a priori* estimation error associated with the predicted reduced HR state. Although, in principle, this quantity corresponds to the covariance of the difference between the true state and its prediction, the true state $\boldsymbol{\psi}_k^{\text{HR}}$ is unknown during online estimation. Therefore, $\mathbf{P}_{k|k-1}$ is computed recursively by propagating the previous updated covariance through the state-transition model and adding the process noise covariance, namely

$$\mathbf{P}_{k|k-1} = \mathbf{F}\,\mathbf{P}_{k-1|k-1}\,\mathbf{F}^{\mathsf{T}} + \mathbf{Q}\,. \tag{18}$$

The corrected estimate can then be written compactly as

$$\hat{\boldsymbol{\psi}}_{k|k}^{\text{HR}} = \mathbf{F}\,\hat{\boldsymbol{\psi}}_{k-1|k-1}^{\text{HR}} + \mathbf{K}_k\left(\boldsymbol{\psi}_k^{\text{LR}} - \mathbf{C}\,\mathbf{F}\,\hat{\boldsymbol{\psi}}_{k-1|k-1}^{\text{HR}}\right)\,, \tag{19}$$

which highlights that the updated HR state is obtained by correcting the model-based prediction through the mismatch between the measured LR coefficients and the corresponding projection of the predicted HR state from the model. The associated covariance matrices are updated according to the standard Kalman filter recursion (Kalman 1960; Welch and Bishop 1995). In the present implementation, the initial state is obtained by applying the pseudoinverse of $\mathbf{C}$ to the first available LR measurement.

2.2.3 Method II—LSE-based estimation with Kalman filtering

The second approach decouples the spatial mapping from the temporal regularization. A LSE is first used to map the LR coefficient vector directly onto an instantaneous estimate of the reduced HR state. A Kalman filter is then applied only as a temporal denoising and regularization stage. This method is hereinafter referred to as *LSE*.

The LSE operator $\mathbf{M} \in \mathbb{R}^{r\times n}$ is identified offline from the training data by ridge-regularized least squares,

$$\mathbf{M} = \mathbf{A}_{\text{train}}^{\text{HR}}\mathbf{A}_{\text{train}}^{\text{LR}^{\mathsf{T}}}\left(\mathbf{A}_{\text{train}}^{\text{LR}}\mathbf{A}_{\text{train}}^{\text{LR}^{\mathsf{T}}} + \lambda\mathbf{I}_n\right)^{-1} \tag{20}$$

where λ is the same regularization parameter introduced in Eq. (12) and set also in this case to $\lambda = 10^{-12}$. In contrast to $\mathbf{C}$ in Eq. (12), which is used in Method I as a forward operator to predict the LR coefficients associated with a given reduced HR state, $\mathbf{M}$ directly provides a statistical LR-to-HR mapping and is therefore used here to reconstruct an instantaneous estimate of the reduced HR coefficients from the observed LR ones.

At each time step, the instantaneous LSE estimate of the reduced HR coefficients is obtained as

$$\tilde{\boldsymbol{\psi}}_k^{\text{HR}} = \mathbf{M}\,\boldsymbol{\psi}_k^{\text{LR}}\,. \tag{21}$$

This estimate is optimal in the regularized least-squares sense over the training data, but it remains affected by frame-to-frame noise inherited from the LR measurement and tends to attenuate the variance of the least observable higher-order modes.

The Kalman filter is then introduced only as a temporal filtering stage acting on the LSE-based estimate. The reduced HR state is assumed to evolve according to

$$\boldsymbol{\psi}_k^{\text{HR}} = \mathbf{F}\,\boldsymbol{\psi}_{k-1}^{\text{HR}} + \mathbf{w}_k\,, \tag{22}$$

while the measurement model is written directly in the HR coefficient space as

$$\tilde{\boldsymbol{\psi}}_k^{\text{HR}} = \boldsymbol{\psi}_k^{\text{HR}} + \boldsymbol{\eta}_k\,, \tag{23}$$

where $\boldsymbol{\eta}_k$ is a zero-mean noise term with covariance $\mathbf{R}_{\mathrm{LSE}}$.

Because the LSE output already provides an estimate in the reduced HR space, no additional measurement operator that provides a mapping between the LR and HR space is needed, and the Kalman filter acts directly on the mismatch between the model prediction and the LSE-based estimate. The Kalman gain, therefore, simplifies to

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\left(\mathbf{P}_{k|k-1} + \mathbf{R}_{\mathrm{LSE}}\right)^{-1}, \tag{24}$$

where $\mathbf{P}_{k|k-1}$ denotes the predicted state-error covariance matrix. The corrected estimate can then be written as

$$\hat{\boldsymbol{\psi}}^{\mathrm{HR}}_{k|k} = \mathbf{F}\,\hat{\boldsymbol{\psi}}^{\mathrm{HR}}_{k-1|k-1} + \mathbf{K}_k\left(\mathbf{M}\,\boldsymbol{\psi}^{\mathrm{LR}}_k - \mathbf{F}\,\hat{\boldsymbol{\psi}}^{\mathrm{HR}}_{k-1|k-1}\right), \tag{25}$$

which highlights that the updated state is obtained by correcting the model prediction through the discrepancy between the instantaneous LSE estimate and the predicted reduced HR state. The associated covariance matrices are updated according to the standard Kalman filter recursion.

The process noise covariance is again estimated from Eq. (14), whereas the measurement noise covariance is estimated from the error between the LSE estimate and the reference HR coefficients on the validation set,

$$\begin{aligned}\mathbf{R}_{\mathrm{LSE}} &= \frac{1}{N_{\mathrm{val}}-1}\mathbf{E}_{\mathrm{LSE}}\mathbf{E}^{\mathsf{T}}_{\mathrm{LSE}},\\ \mathbf{E}_{\mathrm{LSE}} &= \tilde{\mathbf{A}}^{\mathrm{HR}}_{\mathrm{val}} - \mathbf{A}^{\mathrm{HR}}_{\mathrm{val}}.\end{aligned} \tag{26}$$

2.2.4 Method III—Variance-rescaled LSE-based estimation

The third method extends Method II by introducing a heuristic, mode-wise variance correction during the estimation stage. Its purpose is to compensate for the systematic attenuation of the least observable higher-order HR modes, which is a known limitation of regression-based estimators (Ewing and Citriniti 1999; Tinney et al. 2006). The correction is identified offline on the training set and then applied during the online stage. This method is hereinafter referred to as *LSE+VR*.

A diagonal rescaling matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{r\times r}$ is introduced as

$$\boldsymbol{\Gamma} = \mathrm{diag}\left(\left(\mathbf{A}^{\mathrm{HR}}_{\mathrm{train}}\mathbf{A}^{\mathrm{HR}^{\mathsf{T}}}_{\mathrm{train}}\right) \oslash \left(\tilde{\mathbf{A}}^{\mathrm{HR}}_{\mathrm{train}}\tilde{\mathbf{A}}^{\mathrm{HR}^{\mathsf{T}}}_{\mathrm{train}}\right)\right)^{\circ\frac{1}{2}}, \tag{27}$$

where $\mathrm{diag}(\cdot)$ retains only the diagonal entries of its argument as a diagonal matrix, $\oslash$ indicates the Hadamard elementwise division, and the $\circ\frac{1}{2}$ superscript indicates the elementwise Hadamard square root. Each diagonal entry of $\boldsymbol{\Gamma}$ then equals the ratio of the empirical standard deviation of the corresponding true HR coefficients to that of its LSE estimate over the training set: multiplying the prediction by $\boldsymbol{\Gamma}$ restores the correct variance level on average, hence the name *variance rescaling*.

Accordingly, the intermediate estimate entering the Kalman correction is redefined as

$$\tilde{\boldsymbol{\psi}}^{\mathrm{HR}}_k = \boldsymbol{\Gamma}\,\mathbf{M}\,\boldsymbol{\psi}^{\mathrm{LR}}_k. \tag{28}$$

Method III therefore retains the same dynamical model, Kalman gain, and state-update equations introduced in Method II, the only difference being the use of the variance-rescaled intermediate estimate in place of the standard LSE prediction.

The process noise covariance is again estimated from Eq. (14), whereas the measurement noise covariance is obtained from the error between the variance-rescaled estimates and the reference HR coefficients on the validation set,

$$\begin{aligned}\mathbf{R}_{\mathrm{VR}} &= \frac{1}{N_{\mathrm{val}}-1}\mathbf{E}_{\mathrm{VR}}\mathbf{E}^{\mathsf{T}}_{\mathrm{VR}},\\ \mathbf{E}_{\mathrm{VR}} &= \tilde{\mathbf{A}}^{\mathrm{HR}}_{\mathrm{val}} - \mathbf{A}^{\mathrm{HR}}_{\mathrm{val}}.\end{aligned} \tag{29}$$

2.3 HR field reconstruction and baseline comparison

For all three methods, once the Kalman-filtered estimate $\hat{\boldsymbol{\psi}}^{\mathrm{HR}}_{k|k}$ of the reduced HR coefficient vector is available, the corresponding fluctuating HR snapshot is reconstructed through the retained HR POD basis as

$$\hat{\mathbf{z}}^{\mathrm{HR}}_k = \boldsymbol{\Phi}^{\mathrm{HR}}_r\boldsymbol{\Sigma}^{\mathrm{HR}}_r\hat{\boldsymbol{\psi}}^{\mathrm{HR}}_{k|k}, \tag{30}$$

Throughout the paper, the tilde accent $(\tilde{\cdot})$ denotes instantaneous (pre-filtering) estimates, such as the LSE output of Eq. (21), whereas the hat accent $(\hat{\cdot})$ denotes the Kalman-filtered estimates; the field reconstruction is always based on the latter. The full HR velocity field is then recovered by reshaping $\hat{\mathbf{z}}^{\mathrm{HR}}_k$ onto the HR grid and adding $\overline{\mathbf{U}}^{\mathrm{HR}}$.

The proposed methods are compared against direct cubic interpolation of the LR velocity field onto the HR grid. The choice is driven by the need to perform the spatial upsampling in real time, thus preventing the use of more advanced super-resolution or data-assimilation strategies. Unlike the reduced-order estimators introduced above, cubic interpolation does not exploit either the temporal evolution of the flow or the cross-resolution relationships learned from the paired HR/LR training data. It is therefore adopted as a reference baseline to assess the benefit of the proposed data-driven approaches.

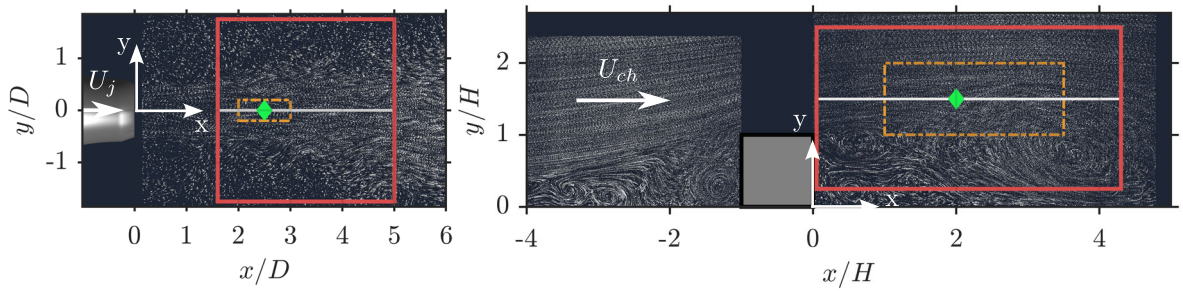


Fig. 2 Representative multi-exposed EBV pseudo-snapshots of the two experimental configurations considered in the present work: submerged round jet (left) and channel flow over a square rib (right). The images are formed by accumulating four consecutive pulses for the jet and ten consecutive pulses for the rib case, in order to enhance the visualization of the flow structures. The so-constructed images are not used for the data processing described in Sec 3.4. The red rectangle marks the resolved region retained for the present study; the orange dashed rectangle identifies the subdomain used for the temporal premultiplied PSD analysis; the green marker denotes the probe location used for the autocorrelation analysis; the white horizontal line indicates the sampling line adopted for the spatial PSD. For detailed sketches of the full experimental setups, the reader is referred to Franceschelli et al. (2025)

3 Experimental datasets and processing pipeline

The estimation framework introduced in Sect. 2 is evaluated on two experimental turbulent flow datasets acquired via pulsed-EBIV: a submerged axisymmetric water jet and an air channel flow obstructed by a spanwise square rib. These two cases were selected as representative of broadband flows, in which energy is distributed over a wide range of spatial and temporal scales. This makes them a challenging testbed for reduced-order estimation and cross-resolution reconstruction.

The pulsed-EBIV principle adopted here follows the approach introduced by Willert (2023). The channel-flow dataset is based on the experimental campaigns documented in Franceschelli et al. (2025), while the jet-flow one is obtained from an upgraded facility compared to the one described in the same work, but with similar optical arrangement. Both the setups are summarized here for completeness, together with the specific processing choices adopted in the present work. In both configurations, the coordinate system shown in Fig. 2 is adopted, with x and y denoting the streamwise and transverse directions, respectively. The main parameters of the two experimental configurations and of the corresponding LR and HR processing pipelines are summarized in Table 1.

Table 1 Summary of experimental and processing parameters for the two flow configurations

Parameter	Jet flow	Channel flow
Medium	Water	Air
Characteristic length	$D = 30$ mm	$H = 8.18$ mm
Reference velocity U_{ref}	0.2 m/s	2.8 m/s
Reynolds number	$Re_D \approx 6\,000$	$Re_H \approx 1\,500$
Acquisition rate	100 Hz	5 kHz
Laser pulse width	1 ms	< 50 ns
EBV camera	Prophesee EVK4 (Sony IMX636)	
Spatial resolution [pixel/mm]	≈ 6.5	≈ 30.3
Pixel displacement corresponding to U_{ref} [pixel/pseudo-frame]	≈ 13.0	≈ 17.0
LR processing (real-time pipeline)		
Algorithm	Single-pass cross-correlation	
IW size [pixel]	48×48	64×64
Overlap	0%	
Window deformation	No	
HR processing (offline reference)		
Algorithm	Multi-frame, multi-pass cross-correlation	
Temporal stencil	3 frames	5 frames
IW size [pixel]	32×32	64×64
Overlap	75%	
Window deformation	Yes	
HR/LR vector/spacing ratio	6	4
Total snapshots N_s	9 000	

3.1 Submerged water jet

The first dataset corresponds to a submerged round jet issuing horizontally into a water tank of dimensions $900 \times 500 \times 450$ mm^3, corresponding to the tank length, width, and height, aligned with the x-, z-, and y-directions, respectively. The nozzle exit diameter is $D = 30$ mm, and the bulk exit velocity is approximately $U_j = 0.2$ m/s, yielding a Reynolds number of $\text{Re}_D = U_j D/\nu \approx 6000$, with ν being the kinematic viscosity. The tank is seeded with polyamide tracer particles of nominal diameter 56 μm.

The measurement plane, passing through the jet axis, is illuminated by a pulsed laser sheet generated by a low-cost continuous-wave diode laser originally designed for engraving applications (LaserTree LT-40W-AA, with 5 W optical power), operated in pulse-width-modulated mode. The light sheet thickness is below 2 mm at the measurement location, as estimated by imaging the sheet cross section on a calibration target during the optical alignment. The pulse repetition rate is $f_j = 100$ Hz, with a pulse width of 1 ms.

Data are recorded using a Prophesee EVK4 event-based camera equipped with a Sony IMX636 sensor featuring a resolution of 1280×720 pixels and a pixel pitch of 4.86 μm. The optical magnification yields a spatial resolution of approximately 6.5 pixel/mm over a field of view of about 190×110 mm^2. The resolved area considered in the present study spans $1.6 \leq x/D \leq 5$ in the streamwise direction and $-1.75 \leq y/D \leq 1.75$ in the radial direction. The acquired dataset consists of 9000 velocity snapshots.

A representative multi-exposed EBV pseudo-snapshot of the jet configuration is shown in Fig. 2. For visual clarity, the image is formed here by combining four consecutive laser pulses. The figure also reports the spatial domains and sampling locations used in the analyses presented later in this work, including the resolved region, the subdomain adopted for the temporal premultiplied Power Spectral Density (PSD), the probe location used for the autocorrelation analysis, and the sampling line used for the spatial PSD. For a detailed sketch of the complete experimental setup, the reader is referred to Franceschelli et al. (2025).

3.2 Channel flow over a square rib

The channel-flow dataset is derived from the experimental campaign described by Franceschelli et al. (2025), to which the reader is referred for full details of the facility, optical arrangement, seeding system, and camera calibration procedure. Briefly, a spanwise square rib of height $H = 8.18$ mm is mounted on the lower wall of a square-section channel of width $W = 76$ mm. The free-stream velocity is $U_{ch} \approx 2.8$ m/s, corresponding to a rib Reynolds number of $\mathrm{Re}_H \approx 1500$. The flow is seeded with paraffin droplets of nominal diameter 1 μm, injected into the upstream settling chamber, and illuminated by a high-speed pulsed Nd:YVO4 laser (Innolas/Iradion Nanio Air 532-10-V-SP). The laser pulse duration is ~ 50 ns and the repetition rate is $f_{ch} = 5$ kHz. EBV data are acquired using the same Prophesee EVK4 camera employed for the jet experiment, with a spatial resolution of 30.3 pixel/mm over a field of view of approximately 42 mm $\times$ 24 mm, corresponding to about $5H \times 3H$.

In the present work, only the wake region downstream of the rib is considered, covering $0 \leq X/H \leq 4.25$ in the streamwise direction and $0.25 \leq Y/H \leq 2.25$ in the wall-normal direction, with the origin located at the intersection between the channel wall and the rear-facing side of the rib. This region includes the separated shear layer, the recirculation bubble, the reattachment zone, and the initial downstream recovery, thus encompassing the most dynamically active portion of the flow.

The complete dataset contains 30 000 velocity snapshots. In the present study, a contiguous subset of 9 000 snapshots is retained to ensure consistency with the jet dataset.

A representative multi-exposed EBV pseudo-snapshot of the rib configuration is also reported in Fig. 2. In this case, the visualization is formed by combining ten consecutive laser pulses in order to better highlight the wake structure. As for the jet case, the figure identifies the resolved domain retained for the present study together with the subdomain and sampling locations used for the spectral and correlation analyses discussed in Sect. 4.

3.3 EBV camera data processing

The raw asynchronous event stream is partitioned into pseudo-snapshots following the pulsed-EBIV procedure introduced by Willert (2023). Each laser pulse generates a burst of positive-polarity events as the illuminated tracer particles cross the per-pixel contrast threshold of the EBV sensor. The laser trigger and the event stream are synchronized through a common TTL signal, allowing the accumulation of all positive events associated with a given laser pulse over a user-defined accumulation time. A pseudo-image is therefore obtained by accumulating all positive events registered within a prescribed temporal window starting from the onset of each laser pulse. The resulting images are binary: pixels that register at least one event are assigned a value of unity, whereas pixels with no activation are assigned zero.

For both experimental cases, the binary pseudo-images are subsequently smoothed with a Gaussian kernel of standard deviation 0.75 pixel prior to cross-correlation processing. In addition, the internal bias settings of the EVK4 camera were adjusted to suppress negative-polarity events, thereby maximizing the available bandwidth for positive events and reducing the risk of arbiter saturation.

3.4 Low- and high-resolution EBIV processing

The same sequence of EBIV pseudo-snapshots is processed in two different ways: a computationally lightweight LR pipeline representative of a realistic real-time implementation, and a more accurate HR pipeline employed offline to generate the reference data used for training and validation.

The LR velocity fields are obtained by single-pass cross-correlation with a fixed interrogation window, without window deformation or iterative refinement. Since only

one correlation step is performed for each window pair, this procedure is representative of the computational budget available in a real-time implementation. No additional post-processing is applied apart from the universal outlier detection (Westerweel and Scarano 2005).

The HR velocity fields are instead computed using an in-house multi-frame, multi-pass cross-correlation algorithm. At each time step, a stencil of consecutive pseudo-snapshots is used to improve the quality of the correlation peak, which is particularly beneficial for EBIV data, where individual pseudo-images are binary and contain limited intensity information. The resulting HR fields provide a spatially refined and temporally regularized reference that the estimator is trained to recover from the corresponding LR measurements.

Table 1 reports the interrogation window sizes, overlaps and temporal stencil dimensions adopted for each dataset in both processing ways.

A remark is due on the different interrogation window overlap adopted in the two pipelines (0% for LR, 75% for HR). Overlapping windows increase the vector density but not, per se, the spatial resolution: cross-correlation over a window of linear size L_{IW} acts as a spatial moving-average filter on the underlying displacement field, with a frequency response set by the window size and by the selected weighting window (Astarita 2007). Velocity fluctuations of streamwise wavelength λ_x are thus increasingly attenuated as λ_x approaches L_{IW} and can be considered well-resolved only for $\lambda_x \gtrsim 2L_{IW}$; moreover, neighbouring HR vectors share particle-image information and exhibit a processing-induced spatial correlation over distances of the order of L_{IW}, regardless of the overlap. For the present HR processing, this bound corresponds to well-resolved content for $D/\lambda_x \lesssim 3$ in the jet case and $H/\lambda_x \lesssim 2$ in the channel-flow case. The actual role of the 75% overlap is to sample this window-filtered field without aliasing: at 0% overlap the sampling spacing equals the filter size itself, so that the attenuated but non-negligible content in the band $L_{IW} \lesssim \lambda_x \lesssim 2L_{IW}$ is aliased, whereas the fourfold oversampling represents it cleanly. The genuine gains of the HR pipeline over the LR one come instead from a combination of the smaller interrogation window (although not strictly required), the window deformation, which improves the spectral response towards the filter bound (Scarano 2002; Astarita 2007), and the multi-frame temporal stencil, which improves the correlation signal-to-noise ratio. Notably, for the channel-flow case the two pipelines share the same window size, so that the HR reference mainly provides denser, alias-free sampling and higher fidelity rather than a wider resolved wavenumber range. The learned LR-to-HR mapping therefore targets the HR representation as produced by the offline pipeline, including its correlation structure: the estimators recover physically independent fine-scale content only up to the effective spatial resolution of the HR processing and reproduce its smoothing characteristics beyond that limit. This should be kept in mind when interpreting the spectral analyses of Sect. 4.4, where the recovered small-scale content is bounded by the resolution of the HR reference itself.

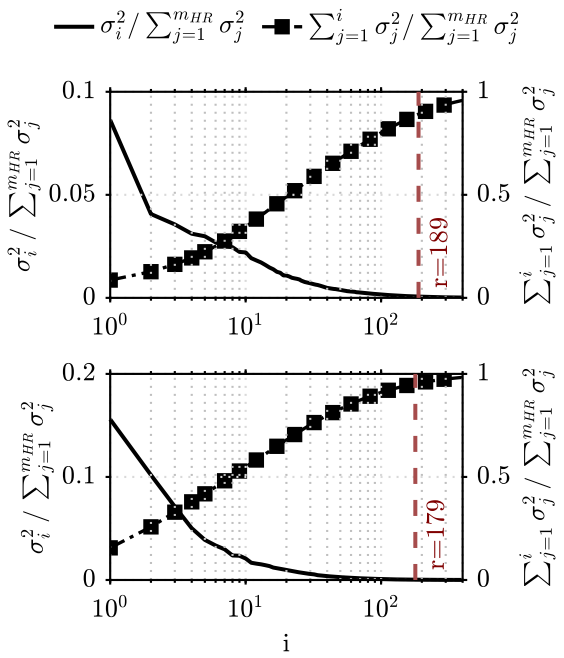


Fig. 3 POD modal energy spectrum and cumulative POD energy for the two considered HR datasets: top, jet flow; bottom, channel flow. The solid black line refers to the left vertical axis and reports the normalized contribution of the i-th mode, $\sigma_i^2 / \sum_{j=1}^{m} \sigma_j^2$. The black dashed line with square markers refers to the right vertical axis and reports the corresponding cumulative contribution, $\sum_{j=1}^{i} \sigma_j^2 / \sum_{j=1}^{m} \sigma_j^2$. The red vertical dashed line indicates the truncation rank r selected by the elbow criterion

The datasets are then divided into training, validation, and testing subsets. The training subset is used to construct the reduced-order spaces and identify the mapping operators, the validation subset is used to estimate the Kalman filter noise statistics (Sect. 2.2.2), and the testing subset is reserved for the final performance assessment. The dataset is split into 4 500 snapshots for training, 1 500 for validation, and 2 000 for testing, with two buffer regions of 500 snapshots inserted between consecutive subsets to reduce temporal correlation.

4 Results

This section evaluates the three proposed estimation strategies on the jet and channel-flow datasets. The discussion is organized hierarchically, from instantaneous field reconstruction to increasingly demanding statistical and spectral diagnostics. Throughout the section, the reference truth used for quality metrics is assumed to be the HR low-order reconstruction (referred in short as HR-LOR), obtained by

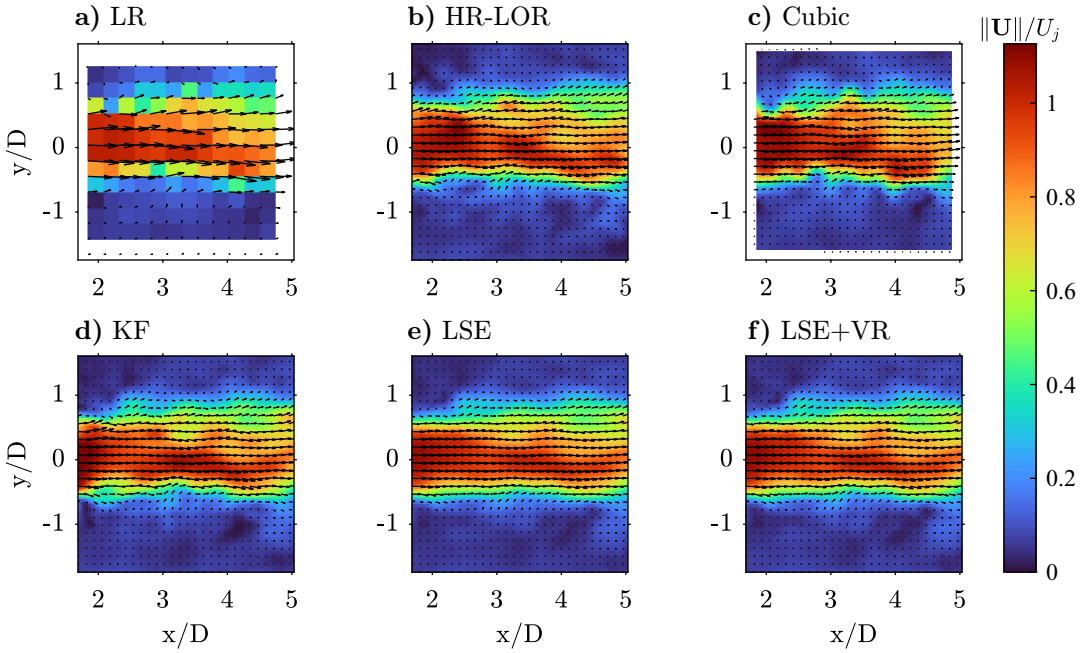


Fig. 4 Comparison of instantaneous velocity field realizations for the jet-flow case. Panels show: **a** input LR field; **b** reference HR-LOR field; **c** HR field obtained by direct cubic interpolation of the LR measurement; **d** HR estimate obtained with KF; **e** HR estimate obtained with LSE; and **f** HR estimate obtained with LSE+VR. Colour contours represent the instantaneous velocity magnitude $\|\mathbf{U}\|$, normalized by the jet bulk velocity U_j, while arrows indicate the in-plane velocity vectors. For readability, one vector every three is shown in the HR panels

projecting the HR dataset onto the truncated POD basis introduced in Sect. 2. The three estimators (KF, LSE, and LSE+VR) are compared against this reference and against direct cubic interpolation of the LR field onto the HR grid, which is adopted as a purely spatial, non-model-based baseline.

It must be stressed that the HR-LOR reference is not an independent ground truth. Rather, it is the low-order projection of the HR fields obtained from the same event stream through the offline processing pipeline. The reported errors therefore quantify how closely the real-time estimators recover the offline-grade reduced-order representation, rather than their absolute deviation from the true flow state. This is consistent with the objective of the present framework, which is precisely to make an offline-grade representation available at real-time rates. The fidelity of this target with respect to conventional velocimetry has, moreover, already been established: in Franceschelli et al. (2025), pulsed-EBIV measurements on the same flow configurations were benchmarked against synchronized, simultaneous PIV measurements—which can fairly be regarded as the reference—with the low-order reconstructions deviating from their PIV counterparts by errors of the order of 0.5% of the peak turbulent kinetic energy for the jet flow, and a few times higher for the channel flow at the truncation ranks of interest here.

4.1 POD modal energy distribution and retained energy

The resulting POD modal energy distributions for the two HR datasets are reported in Fig. 3. Since the energy captured by the i-th POD mode within the dataset equals the square of the corresponding singular value, σ_i^2, the singular value distribution coincides with the POD modal energy spectrum, and its running sum with the cumulative retained energy. For each case, the figure shows the normalized modal energy σ_i^2 together with its cumulative contribution as functions of the mode index i. The red vertical line indicates the truncation rank selected by the *elbow criterion* introduced above. The adopted threshold yields $r = 189$ for the jet case and $r = 179$ for the channel-flow case, corresponding to retained cumulative energies of 89.1% and 95.2%, respectively. In both datasets, the singular value decay is gradual rather than characterized by a sharp cut-off. As a result, the selected truncation lies in the high-order portion of the spectrum while still retaining most of the resolved modal energy. This behaviour highlights the broadband character of the two flow cases and the wide distribution of energy across spatial scales, making these datasets particularly challenging for reduced-order reconstruction and estimation.

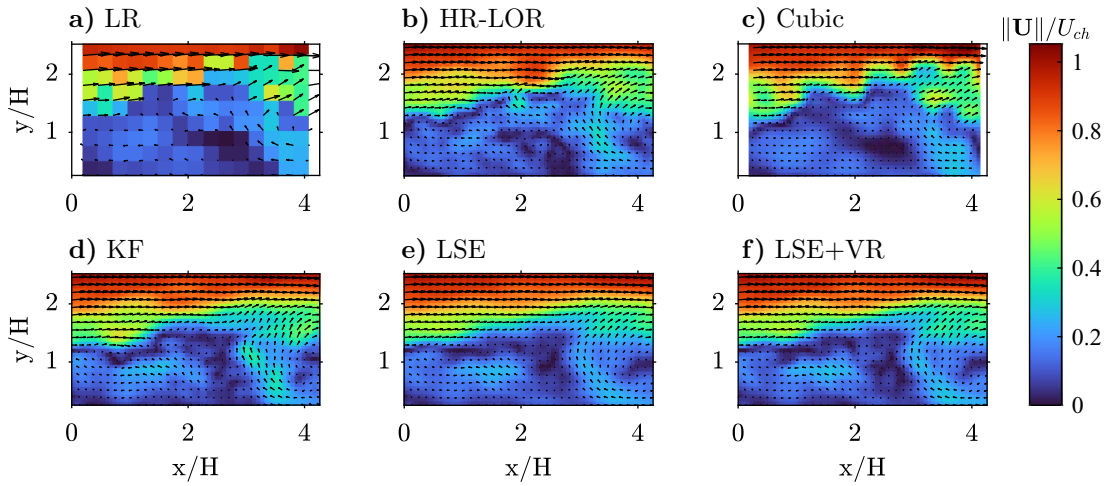


Fig. 5 Comparison of instantaneous velocity field realizations for the channel-flow case. Panels show: **a** input LR field; **b** reference HR-LOR field; **c** HR field obtained by direct cubic interpolation of the LR measurement; **d** HR estimate obtained with KF; **e** HR estimate obtained with LSE; and **f** HR estimate obtained with LSE+VR. Colour contours represent the instantaneous velocity magnitude $\|\mathbf{U}\|$, normalized by the free-stream channel-flow velocity U_{ch}, while arrows indicate the in-plane velocity vectors. For readability, one vector every two is shown in the HR panels

4.2 Instantaneous field reconstruction

A first qualitative assessment is presented in Figs. 4 and 5 for the jet and channel-flow cases, respectively, depicting the instantaneous velocity fields. Each figure compares the input LR field (a), the HR-LOR reference (b), the field obtained by direct cubic interpolation (c), and the three reconstructed fields obtained with KF (d), LSE (e), and LSE+VR (f). The colour contours represent the instantaneous velocity magnitude $\|\mathbf{U}\|$, while the quiver plot indicates the in-plane velocity vector. The LR field provides only a coarse and noisy representation of the instantaneous flow, whereas the HR-LOR field reveals substantially finer spatial detail, including sharp velocity gradients and the coherent structures that characterize each flow: the shear layers and vortical roll-up of the jet core, and the separated region downstream of the square rib in the channel-flow case. All three proposed estimators recover the broad spatial organization of the HR-LOR field with good fidelity. The KF estimate provides a smooth reconstruction that closely follows the large-scale coherent features of the reference. The LSE reconstruction captures the same dominant structures, although the smallest-scale features appear less sharply defined. The LSE+VR reconstruction recovers additional small-scale content while maintaining spatial regularity. In contrast, the cubic-interpolation baseline introduces visible high-frequency oscillations and artificial spatial modulations that are absent in the reference field. These artefacts are particularly evident in the channel-flow case, where the interpolation of the noisy LR measurements produces spurious patterns in regions of strong velocity gradient.

4.3 Statistics of the reconstructed fields

The spatial distribution of the in-plane Turbulent Kinetic Energy (TKE), defined here as $\frac{1}{2}\langle u^2(\mathbf{x},t) + v^2(\mathbf{x},t)\rangle$, is reported in Figs. 6 and 7 for the jet and channel-flow cases, respectively.

The HR-LOR reference reveals clear and physically consistent high-energy regions. In the jet, the TKE is concentrated along the two shear layers, where the strongest fluctuations are associated with the roll-up and downstream convection of coherent structures. In the channel flow, the highest energy levels are found in the separated shear layer and in the downstream recovery region.

The three proposed estimators recover these energetic regions with good overall consistency. Among them, the KF result is the closest to the reference in both spatial footprint and amplitude, with high-energy regions appearing at the correct locations and with a realistic spatial spreading. The LSE reconstruction also captures the overall topology of the TKE field, but the energetic regions appear smoother and less sharply defined, with reduced peak levels, indicating an attenuation of the turbulent fluctuations. This behaviour is consistent with the tendency of linear regression estimators to smooth the fluctuating content, especially for less energetic and less correlated components of the flow. The LSE+VR correction mitigates this effect by restoring part of the missing modal variance, leading to a better recovery of the energy distribution while preserving a physically regular spatial organization.

The cubic-interpolation baseline performs substantially worse in both cases. The reconstructed field exhibits peri-

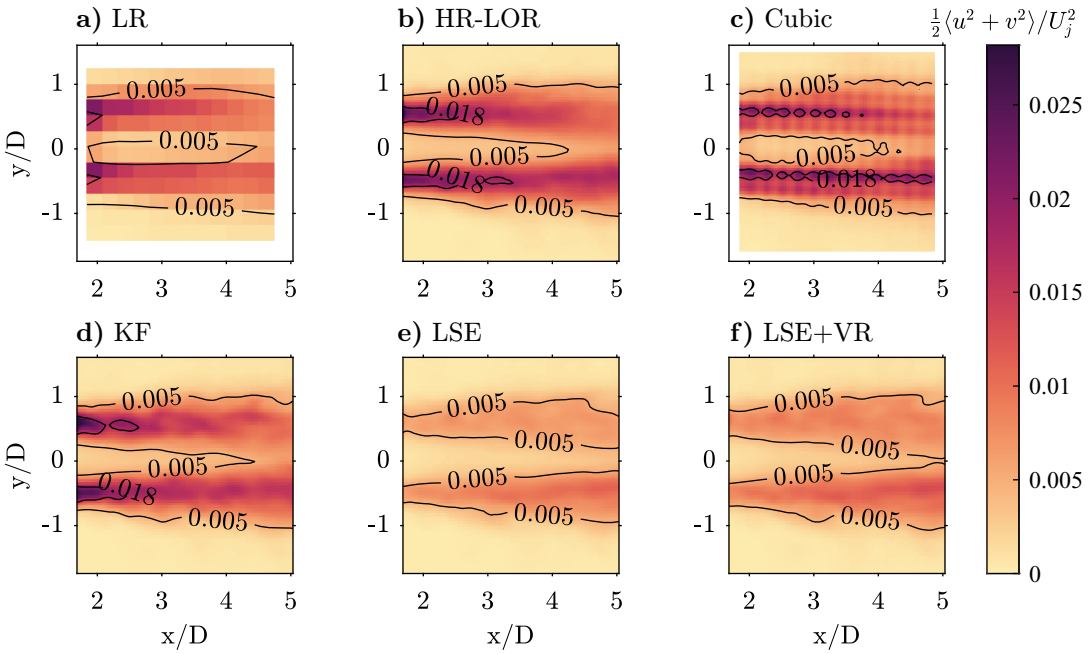


Fig. 6 Spatial distribution of TKE for the jet-flow case, normalized by the square of the jet bulk velocity U_j. Panels show: **a** LR input; **b** HR-LOR reference; **c** cubic interpolation; **d** KF; **e** LSE; and **f** LSE+VR. Black isolines denote selected TKE levels

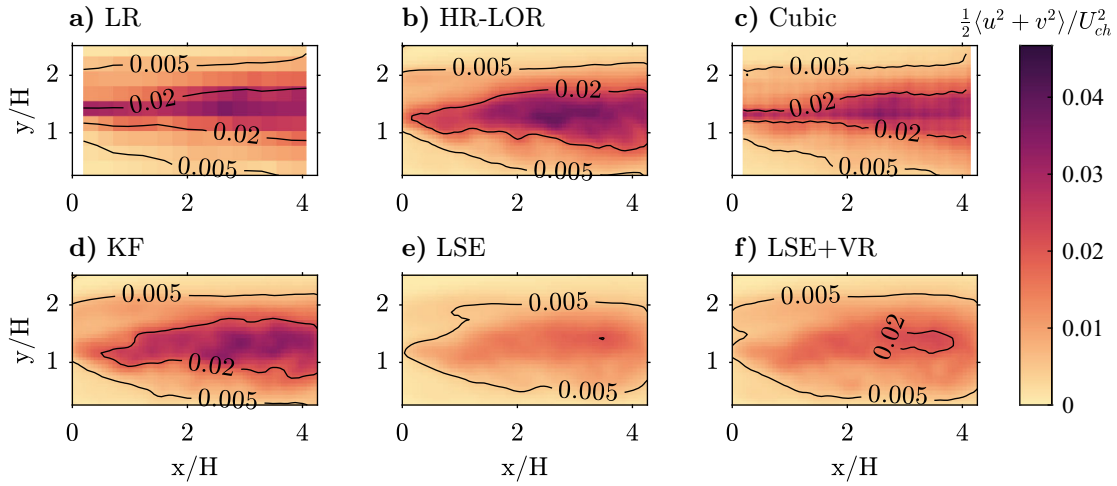


Fig. 7 Spatial distribution of TKE for the channel-flow case, normalized by the square of the free-stream channel-flow velocity U_{ch}. Panels show: **a** LR input; **b** HR-LOR reference; **c** cubic interpolation; **d** KF; **e** LSE; and **f** LSE+VR. Black isolines denote selected TKE levels

odic energy modulations that are absent in the reference, visible as alternating bands or ripples of higher and lower TKE, particularly along the jet shear layers and in the energetic region of the channel flow. These artefacts arise from grid imprinting: when under-resolved and noisy LR fluctuating fields are interpolated onto a finer grid, the interpolation operator transfers the footprint of the original sampling pattern into the reconstructed field. Because TKE is a quadratic quantity, these non-physical oscillations are amplified and become much more evident than in the instantaneous velocity snapshots.

Overall, these results show that the main benefit of the proposed estimators is not limited to reconstructing plausible instantaneous fields, but extends to recovering fluctuation statistics with a physically consistent spatial organization and without the grid-imprinting artefacts observed with interpolation.

4.4 Temporal and spatial spectral content

The spectral fidelity of the reconstructed fields is assessed through pre-multiplied temporal and spatial spectra of the

streamwise velocity fluctuation u. In the following, $\langle\cdot\rangle_T$ denotes temporal averaging and $\langle\cdot\rangle_\Omega$ denotes spatial averaging over the spatial subdomain Ω.

The quantity reported in the temporal spectral analysis of Fig. 8 is the premultiplied temporal PSD of the streamwise velocity fluctuation, $St \cdot \langle S_{uu}\rangle_\Omega$, computed from the local time signals within the subdomain Ω, and the Strouhal number is defined as $St = fD/U_j$ for the jet and $St = fH/U_{ch}$ for the channel flow. The temporal spectra are then averaged over Ω, defined as $2 \le x/D \le 3$ and $-0.2 \le y/D \le 0.2$ for the jet, and $1 \le x/H \le 3.5$ and $1 \le y/H \le 2$ for the channel flow, corresponding, respectively, to the jet core and the highly turbulent region behind the square rib. The corresponding subdomains are indicated by the orange dashed rectangles in Fig. 2.

The jet subdomain covers a comparatively compact region of size $1D \times 0.4D$, centred on the jet core, where the signature of the jet-column coherent dynamics emerges most clearly and the comparison between estimators is less affected by measurement noise. The channel-flow subdomain is deliberately more extended: it samples a fully turbulent broadband region (representative of the behaviour that would also be observed across a turbulent shear layer) and the wider spatial averaging promotes the statistical convergence of the spectra, at the price of smoothing out strictly local contributions. It has been verified that the reported trends are insensitive to this choice: repeating the analysis over a much smaller nested subdomain leaves the relative behaviour of the estimators unchanged, at the cost of noisier spectra; the jet case trends were likewise verified by recomputing the spectra within the shear layer region. The two cases are therefore regarded as complementary: the former probes the recovery of coherent, spectrally localized dynamics, the latter the recovery of broadband turbulent content.

The premultiplied spatial spectra, shown in Fig. 9, are computed along a horizontal sampling line located at $y/D = 0$ for the jet and at $y/H = 1.5$ for the channel flow, as indicated by the grey horizontal line in Fig. 2. In this case, the spectra are represented in premultiplied non-dimensional form as $(L_{ref}/\lambda_x)\,\langle E_{uu}\rangle_T$, where E_{uu} denotes the one-dimensional streamwise spatial spectrum of u, L_{ref} is a reference length (D for the jet and H for the channel flow), and λ_x is the spatial wavelength.

The temporal spectra show that all three proposed estimators capture the dominant frequency content of the flow. In the jet case, the main spectral peak associated with the dominant coherent dynamics is recovered in both position and amplitude. In the channel-flow case, where the spectrum is broader and less sharply peaked, the low-frequency content is likewise captured correctly. The KF estimator provides the closest agreement with the HR-LOR reference over the widest frequency range, including part of the high-frequency tail, capturing higher energy content with less attenuation, consistently with the TKE results. Standard LSE and LSE+VR follow the same overall spectral shape and retain the dominant dynamics, but with a lower overall energy level, consistent with the attenuated fluctuating fields already observed in the TKE maps. The LSE+VR is observed to perform consistently better than the standard counterpart.

The cubic-interpolation baseline again exhibits the least reliable behaviour. Although it reproduces part of the low-frequency content (albeit with noticeable attenuation), it develops a markedly elevated broadband floor at intermediate and high Strouhal numbers. This excess floor is more than one order of magnitude above the reference over part of the spectrum, and in the jet tail it approaches nearly two orders of magnitude. As a result, the physical spectral decay is partially obscured and the identification of dominant peaks becomes less reliable, particularly in the channel-flow case where the broadband contamination masks part of the spectral structure. The spectra of the raw LR fields, also reported in Fig. 8, make the origin of this floor explicit: the LR spectra follow the reference at low Strouhal numbers, but flatten into an elevated broadband floor at intermediate and high frequencies, which constitutes the spectral signature of the measurement noise of the lightweight single-pass processing. The cubic-interpolation baseline essentially inherits this floor, as expected from a purely spatial operator acting on the LR fields, whereas the proposed estimators suppress it by constraining the reconstruction to the learned reduced-order subspace. These observations are consistent with the known spectral response of conventional PIV (Foucaut et al. 2004).

The spatial spectra reinforce the same interpretation from a complementary perspective. The three data-driven estimators reproduce the HR-LOR spatial energy distribution over a broad range of resolved wavelengths, including part of the small-scale content that is absent from the raw LR field. The raw LR spectra, also included in Fig. 9, follow the reference over the largest resolved scales with a slight energy excess ascribable to measurement noise, but terminate at markedly smaller wavenumbers than all the other curves: the coarse LR vector spacing bounds the shortest resolvable wavelength to the corresponding Nyquist limit, $\lambda_{x,\min} = 2\Delta x_{\mathrm{LR}}$, i.e. $D/\lambda_x \approx 2$ for the jet and $H/\lambda_x \approx 2$ for the channel flow. This early cut-off visualizes directly the spatial information gap that the proposed estimators are required to fill. The KF result provides the closest overall match and, especially in the jet case, nearly overlaps the HR-LOR reference over most of the resolved spectral range. The advantage of the VR correction is particularly evident in this representation: compared with standard LSE, LSE+VR recovers part of the missing energy while preserving the correct spectral trend. In contrast, cubic interpolation departs from the reference already at intermediate wavelengths and exhibits characteristic undulations at wavelengths comparable to the original LR vector

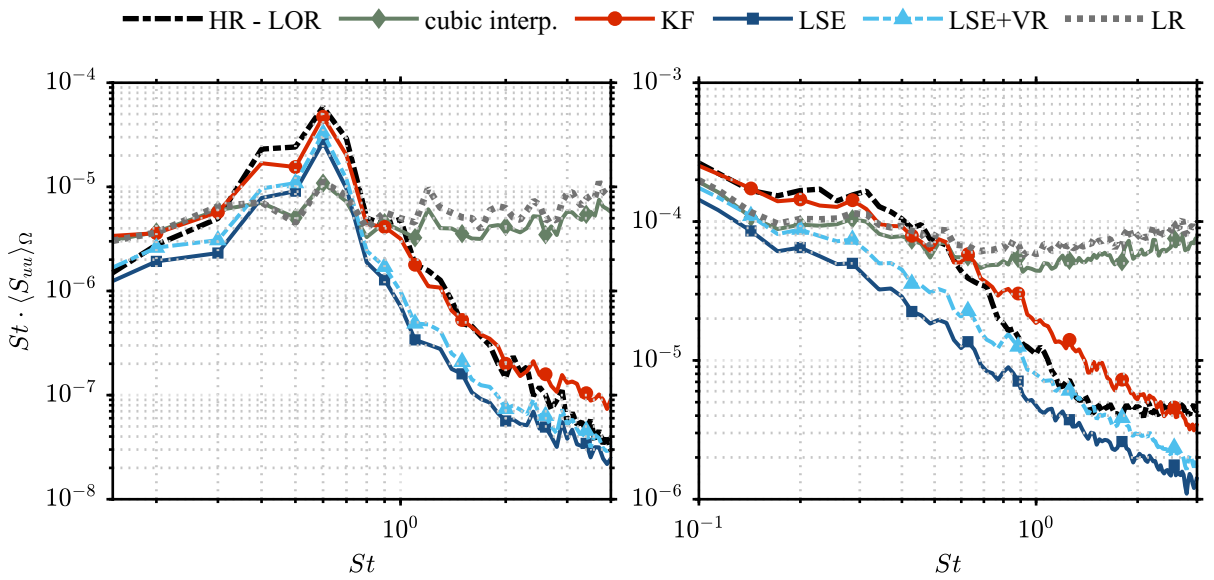


Fig. 8 Premultiplied temporal PSDs of the streamwise velocity fluctuations for the jet (left) and channel-flow (right) cases, normalized, respectively, by U_j and U_{ch}, reported as $St\,\langle S_{uu}\rangle_\Omega$. The spectra are spatially averaged over the subdomain Ω defined by $2 \le x/D \le 3$ and $-0.2 \le y/D \le 0.2$ for the jet, and $1 \le x/H \le 3.5$ and $1 \le y/H \le 2$ for the channel flow, and compared against the HR-LOR reference. The premultiplied spectra are reported in non-dimensional form, with the spectral energy normalized by U_j^2 and U_{ch}^2, respectively. The grey dotted line reports the spectra of the raw LR fields

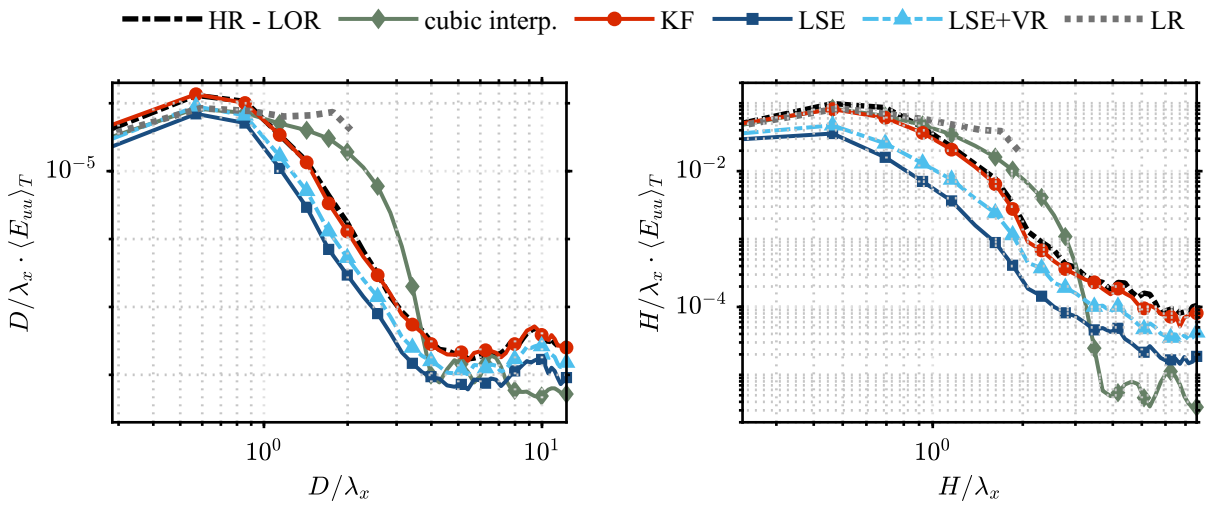


Fig. 9 Premultiplied spatial spectra of the streamwise velocity fluctuations for the jet (left) and channel-flow (right) cases, normalized, respectively, by U_j and U_{ch}, computed along the horizontal lines $y/D = 0$ and $y/H = 1.5$, respectively. The spectra are shown as $(D/\lambda_x)\,\langle E_{uu}\rangle_T$ for the jet and $(H/\lambda_x)\,\langle E_{uu}\rangle_T$ for the channel flow, and compared against the HR-LOR reference. The premultiplied spectra are reported in non-dimensional form, with the spectral energy normalized by U_j^2 and U_{ch}^2, respectively. The grey dotted line reports the spectra of the raw LR fields, whose resolved wavenumber range terminates earlier owing to the coarser LR vector spacing

spacing. In the jet case, this interpretation is quantitatively consistent with the LR sampling pitch. From the image scaling, the LR vector spacing is $\Delta x_{\mathrm{LR}} \approx 7.4$ mm, corresponding to $D/\Delta x_{\mathrm{LR}} \approx 4.1$, namely about four LR vectors across the nozzle diameter. The interpolated spectrum indeed departs from the reference around this scale, displaying an energy deficit (notch) near $D/\lambda_x \approx 4$ followed by a spurious energy excess at slightly smaller wavelengths, $D/\lambda_x \approx 5$–6. These undulations result directly from grid imprinting and confirm that the interpolation baseline does not reconstruct the missing HR content, but instead introduces artificial spatial scales into the field.

From a physical standpoint, these spectra show that the proposed framework not only improves the visual quality of

the reconstructed fields but also preserves the scale distribution of the fluctuating dynamics. This aspect is essential if the reconstructed fields are to be used for reduced-order modelling, system identification, or feedback-oriented state estimation.

4.5 Reduced-order dynamics

To assess the estimation accuracy directly in the reduced space, Figs. 10 and 11 compare the time histories of four selected reduced-order coefficients together with the premultiplied spectrum of a representative one. The left-hand side of each figure reports the temporal evolution of the coefficients associated with modes 1, 6, 25, and 85 over a window of five convective time units. For each mode i, the plotted quantity is the corresponding scalar component of the reduced state, namely $\psi_{k,i}^{\mathrm{HR}}$ for the HR-LOR reference and $\hat{\psi}_{k|k,i}^{\mathrm{HR}}$ for the estimate, normalized by $\sqrt{N_t}$.

The convective time is defined as $t_c = tU_{ref}/L_{ref}$, with $U_{ref} = U_j$ and $L_{ref} = D$ for the jet, and $U_{ref} = U_{ch}$ and $L_{ref} = H$ for the channel flow. The four modes are selected to span the entire retained spectrum: the most energetic mode (mode 1), an intermediate-energy mode still dominated by coherent dynamics (mode-6), and two progressively higher-order, low-energy modes (modes 25 and 85), for which the estimation problem becomes increasingly challenging. The right-hand side shows the premultiplied spectrum of the coefficient associated with mode 6, adopted as representative example because it combines a significant energy content with a well-defined spectral signature, which makes the differences between estimators clearly visible. For visual clarity, only the LSE+VR result is shown among the LSE-based estimators, as in this representation it largely overlaps with the standard LSE result.

This analysis is particularly informative because the reduced coefficients are the actual state variables estimated by the proposed methods. Accurate coefficient reconstruction therefore implies not only a faithful reconstruction of the velocity field, but also a reliable low-dimensional representation of the underlying dynamics.

For the leading modes, i.e. the most energetic modes retrieved at the lowest mode indices, all reconstruction methods recover the overall trend of the HR-LOR reference reasonably well. The KF and LSE+VR estimates are generally the most consistent, reproducing the dominant low-order dynamics with good agreement in both amplitude and phase. Cubic interpolation often follows the large-scale trend as well, but already at low mode number, it exhibits visibly noisier coefficient traces, with superposed fluctuations that are absent from the reference. This behaviour is particularly evident in the channel-flow case, which is more strongly affected by measurement noise. As the mode rank increases, the reconstruction problem becomes progressively harder because the higher-order coefficients carry less energy and lie closer to the noise floor of the LR measurements. In this regime, the interpolation-based estimates become much less trustworthy and may develop large spurious excursions, whereas the model-based estimators retain a meaningful correlation with the reference.

The attenuation of the LSE-based reconstruction remains visible at higher mode number, as illustrated for instance by mode 85 in the channel-flow case. The VR correction partially compensates for this loss by restoring part of the missing modal variance, yielding a reconstruction with amplitudes comparable with KF. The coefficient spectra confirm the same interpretation. In the jet case, KF reproduces the reference spectrum most closely over the widest Strouhal range. In the channel-flow case, the situation partially reverses: the KF spectrum exceeds the reference at high frequencies, indicating that part of the measurement noise is transferred to the estimated coefficients through the correction step, whereas the smoother LSE+VR estimate remains in part closer to the reference over that range. This interchange in the relative performance of the two estimators between the two flow cases reflects the different role played by the linear dynamical model: in the jet, the quasi-periodic coherent dynamics are well captured by the one-step transition model, which allows the KF to filter the measurement noise effectively; in the broadband channel flow, the linear model is less informative, and the KF correction step consequently admits a larger share of the measurement noise. The same trend is observed across the other retained modes, with the KF high-frequency excess visible in the channel-flow case over a wide range of mode indices, and the LSE+VR attenuation persisting at high mode index in both flow cases. It is worth remarking that both estimators consistently reproduce the *shape* of the coefficient spectra (the spectral decay and the distribution of the dominant content) with the discrepancies essentially confined to the energy level, as in the LSE+VR attenuation. By contrast, the spectra of the cubic-interpolation baseline are nearly flat beyond the low-frequency range: the physical spectral decay is lost altogether, indicating that the interpolated coefficients are dominated by noise rather than by the underlying modal dynamics. By contrast, cubic interpolation displays a clearly elevated broadband floor, confirming that it does not simply lose information, but also injects spurious modal activity.

4.6 Temporal coherence from virtual probes

A complementary assessment of the reconstruction quality is provided by a lag-correlation analysis at a virtual probe located at $(x/D, y/D) = (2, 0)$ for the jet and $(x/H, y/H) = (2, 1.5)$ for the channel flow, respectively. For the HR-LOR reference, the temporal autocorrelation of

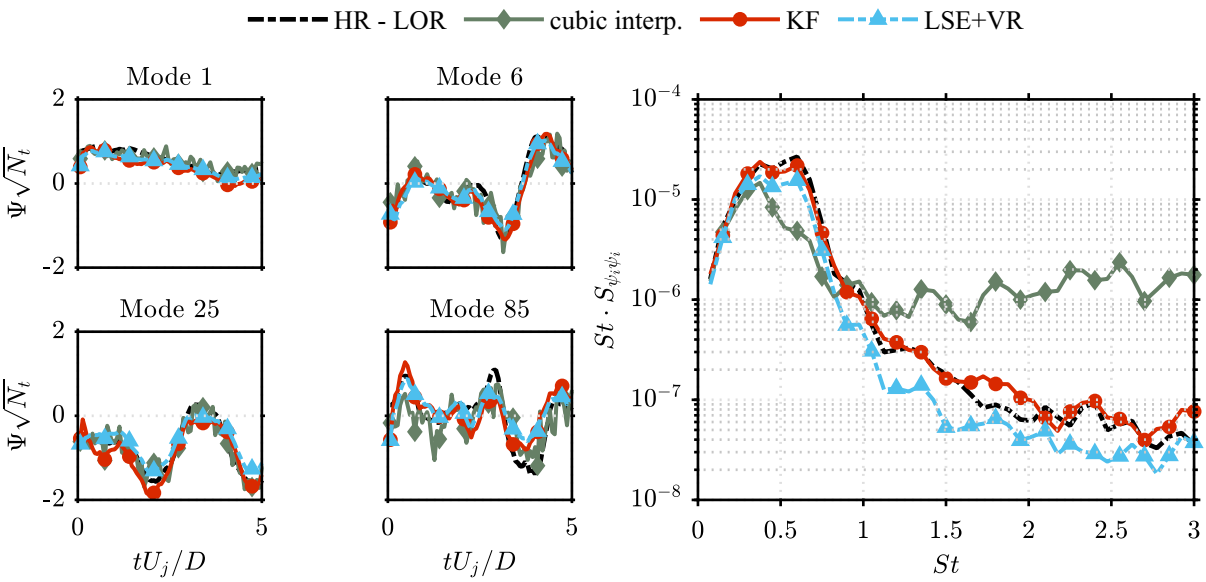


Fig. 10 Jet-flow reduced-order dynamics. Left: time histories of selected POD coefficients (modes 1, 6, 25, and 85) over five convective time unit. Right: premultiplied spectrum of the mode-6 coefficient. Coefficients are normalized as $\psi_i/\sqrt{N_t}$. For visual clarity, only the LSE+VR result is shown among the LSE-based estimators

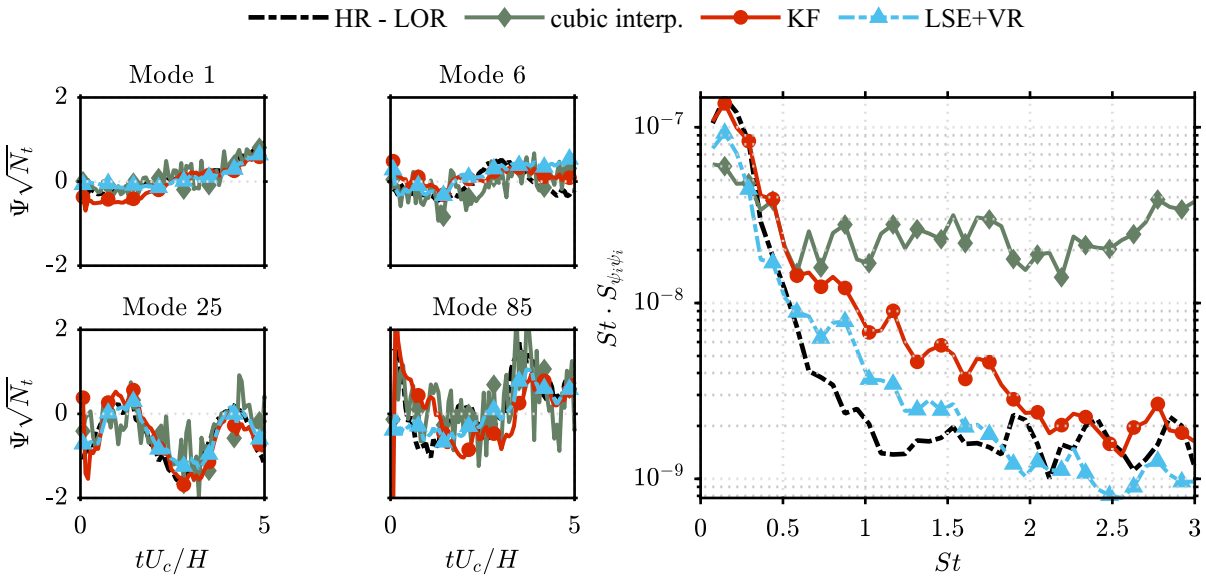


Fig. 11 Channel-flow reduced-order dynamics. Left: time histories of selected POD coefficients (modes 1, 6, 25, and 85) over five convective time units. Right: premultiplied spectrum of the mode-6 coefficient. Coefficients are normalized as $\psi_i/\sqrt{N_t}$. For visual clarity, only the LSE+VR result is shown among the LSE-based estimators

the streamwise velocity fluctuation is reported,

$$R_{uu}(\tau) = \frac{\langle u(t)\, u(t+\tau)\rangle}{\langle u^2(t)\rangle}, \tag{31}$$

which is symmetric by construction. For the estimators, the normalized cross-correlation between the estimated and the reference probe signals is shown instead,

$$R_{\hat{u}u}(\tau) = \frac{\langle \hat{u}(t)\, u(t+\tau)\rangle}{\sigma_{\hat{u}}\, \sigma_u}, \tag{32}$$

where $\sigma_{\hat{u}}$ and σ_u denote the standard deviations of the estimated and reference signals. This cross-statistic is adopted deliberately, since it is the only one able to reveal a temporal misalignment between reconstruction and reference: an artificial time lag would displace the correlation peak away from $\tau = 0$. The results are reported in Fig. 12. For visual clar-

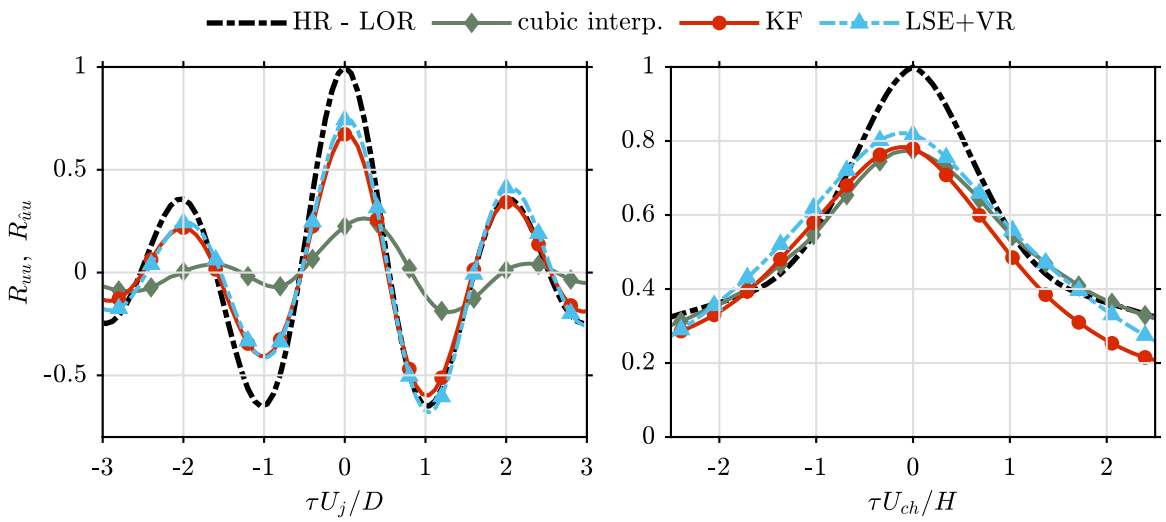


Fig. 12 Lag-correlation analysis of the streamwise velocity fluctuation u from virtual probes located at $(x/D, y/D) = (2, 0)$ for the jet (left) and $(x/H, y/H) = (2, 1.5)$ for the channel flow (right): temporal autocorrelation R_{uu} for the HR-LOR reference, and normalized cross-correlation with the reference, $R_{\hat{u}u}$, for the estimators. The time lag is normalized as $\tau U_j/D$ for the jet and $\tau U_{ch}/H$ for the channel flow. For visual clarity, only the LSE+VR result is shown among the LSE-based estimators

ity, only the LSE+VR result is shown among the LSE-based estimators.

In the jet case, the HR-LOR reference displays the expected quasi-periodic symmetric autocorrelation signature associated with the passage of coherent vortical structures and with a well-defined characteristic timescale. Both KF and LSE+VR reproduce the delay structure well, without introducing an appreciable phase shift. The slight asymmetry of the reconstructed curves with respect to zero lag is an expected feature of this cross-statistic: the estimators are causal, operating on current and past measurements only, so their output does not correlate symmetrically with the past and the future of the reference signal. The correlation peak remains located at $\tau = 0$, confirming that no appreciable artificial time lag is introduced. In terms of correlation amplitude, LSE+VR performs comparably to KF, following the reference more closely over part of the lag range. By contrast, cubic interpolation yields a markedly damped and distorted autocorrelation shape, indicating that the interpolation procedure alters the underlying dynamics.

In the channel-flow case, the autocorrelation decays more smoothly and monotonically, consistent with the broader-band character of the separated-flow dynamics. The differences between the reconstruction methods are smaller than in the jet case. LSE+VR provides the best agreement in terms of correlation amplitude over much of the lag range, while cubic interpolation reproduces the overall correlation shape reasonably well. In this representation, the KF result does not exhibit a clear advantage over interpolation. While the Kalman filter enforces physical consistency, its reliance on a linear dynamical model leads to over-smoothing in the stochastic channel flow; here, local interpolation better captures the high-frequency fluctuations, despite lacking a physics-based formulation.

Overall, these results confirm that the data-driven estimators preserve the temporal organization of the reconstructed signal without introducing an evident artificial time lag, whereas cubic interpolation degrades the correlation shape and therefore distorts the underlying temporal dynamics.

4.7 Global reconstruction error

Finally, a compact scalar measure of reconstruction accuracy is reported in Table 2. For the n-th snapshot and the i-th sampling point of the HR grid, the local reconstruction error is defined as

$$e_{i,n} = \left\| \hat{\mathbf{u}}_n(\mathbf{x}_i) - \mathbf{u}_n^{\mathrm{LOR}}(\mathbf{x}_i) \right\|, \tag{33}$$

where $\mathbf{x}_i$ denotes the location of the i-th HR sampling point. Hatted quantities denote the estimated fields, while the superscript LOR refers to the HR-LOR reference used throughout the paper. The overall error is then defined as the average normalized Root-Mean Squared Error (RMSE) of this local velocity error over the testing dataset:

$$\delta = \frac{1}{N_s\, U_{\mathrm{ref}}} \sum_{n=1}^{N_s} \sqrt{\frac{1}{N_p^{\mathrm{HR}}} \sum_{i=1}^{N_p^{\mathrm{HR}}} e_{i,n}^2}, \tag{34}$$

where N_s is the number of testing snapshots and U_{ref} denotes the reference velocity, namely U_j for the jet and U_{ch} for the channel-flow case.

The results reported in Table 2 show that, over the full resolved domain, all three proposed estimators outperform the cubic-interpolation baseline in both configurations. Among the proposed methods, LSE provides the lowest overall error in both the jet and channel-flow cases, with LSE+VR yielding nearly identical values. The improvement over cubic interpolation is appreciable, although still moderate in absolute terms. This should be interpreted in light of the intrinsic limitations of a pointwise scalar metric. By compressing all spatial and temporal information into a single number, δ cannot discriminate between physically consistent reconstructions and fields contaminated by non-physical artefacts. The metric δ should accordingly be regarded as a complementary summary indicator rather than the primary criterion for assessing reconstruction quality.

The apparent disparity between the field- and spectrum-based diagnostics, where KF often provides the closest agreement with the reference, and the global error, where LSE prevails, can be interpreted in relation to the nature of the estimators. The LSE operator is identified by minimizing the mean-squared error over the training data and is therefore near-optimal by construction with respect to δ-type metrics; this optimality is achieved at the cost of a systematic attenuation of the fluctuations. The KF, by contrast, trades pointwise accuracy for dynamic and spectral consistency, which is rewarded by the energy-based diagnostics but not by a pointwise scalar error. Additionally, to assess the sensitivity of δ to the choice of the spatial domain, the metric has also been evaluated restricting the spatial average to the most energetic regions of the flow, defined as the HR grid points where the reference TKE exceeds a fraction α of its maximum; the resulting values δ_α, together with the corresponding mask coverages, are reported in Table 2 for $\alpha = 0.25$ and $\alpha = 0.5$. The errors increase for all methods, as expected since the metric remains normalized by U_{ref} while the local fluctuation intensity grows, and the LSE-based estimators remain the most accurate in the restricted domains as well, with LSE+VR marginally improving over standard LSE in the channel-flow case. The contrast between the model-based estimators and the interpolation baseline, however, narrows in the energetic regions and, for the strictest threshold in the jet case ($\alpha = 0.5$, isolating the thin high-TKE bands of the shear layers over 8% of the resolved points), the KF pointwise error becomes slightly larger than that of interpolation. This behaviour reflects the different objective of the estimators discussed above: the KF privileges dynamic and spectral consistency over pointwise optimality, retaining fluctuation amplitudes that a pointwise metric penalizes wherever energetic structures are slightly misplaced, whereas smoother estimates are favoured by δ even when they underrepresent the flow dynamics. The domain-restricted analysis thus provides quantitative confirmation of the complementary character of δ-type indicators noted above.

Finally, it is useful to quantify the error introduced by the low-order truncation alone to partially decouple the estimation error from the fidelity of the reference itself. Evaluated with the same metric of Eq. (34), but comparing the full HR fields with their HR-LOR projection over the testing dataset, the truncation error amounts to $\delta_{\text{LOR}} = 0.0550$ and 0.0540 for the jet and channel-flow case, respectively. These values represent lower bounds for the total deviation of any of the proposed estimators from the full HR fields and should be kept in mind when interpreting the errors of Table 2, which are computed with respect to the HR-LOR reference. More precisely, since the estimated fields lie in the retained subspace by construction, while the truncation residual is orthogonal to it, the deviation of the proposed estimators from the full HR fields is the root-sum-square of the two contributions, $\sqrt{\delta^2 + \delta_{\text{LOR}}^2}$, amounting for instance to ≈ 0.088 and ≈ 0.117 for LSE in the jet and channel-flow cases, respectively; this decomposition does not apply to the interpolation baseline, whose output is not confined to the retained subspace. Since the truncation rank is selected by the *elbow criterion* precisely so that the discarded modes are predominantly noise-dominated (Sect. 2.1), the HR-LOR reference acts as a denoised version of the HR fields, and δ_{LOR} partially reflects removed measurement noise; it should therefore be regarded as a conservative upper bound on the information excluded by the reference.

4.8 Computational cost of the reduced-order refinement

To quantify the online overhead introduced by the proposed LR-to-HR refinement stage, a dedicated microbenchmark implemented in C++ was performed on the jet case. The benchmark was implemented following the online workflow of Method I (direct Kalman filter estimation), including the learned state-transition matrix $\mathbf{F}$. Although the exact latent-state update depends on the chosen estimator, the benchmark is intended to provide a representative order-of-magnitude assessment of the additional computational overhead of the reduced-order refinement. The benchmark was carried out on a workstation running Microsoft Windows 11 Pro (build 26200), equipped with an AMD Ryzen Threadripper 7960X CPU (24 physical cores, 48 logical processors) and 256 GB of RAM. All reported timings correspond to single-threaded execution. The benchmark accounts only for the additional algebraic operations required once the LR velocity field is already available, namely: projection of the incoming LR field onto the reduced space, estimation of the HR latent state, and reconstruction of the HR field from the estimated latent coefficients.

The results are reported in Fig. 13. For retained ranks up to $r = 200$, the total added cost remains below 0.5 ms on

Table 2 Normalized RMSE δ (Eq. 34) for the two flow cases, computed over the full resolved domain and restricted to the most energetic regions, δ_α, defined as the HR grid points where the reference TKE exceeds a fraction α of its maximum (the spatial average of Eq. 34 being restricted to those points). The last row reports the fraction of resolved points retained by each mask. The best result in each column is highlighted in bold

	Jet flow			Channel flow		
Method	δ	$\delta_{0.25}$	$\delta_{0.50}$	δ	$\delta_{0.25}$	$\delta_{0.50}$
Cubic interpolation	0.0876	0.1235	0.1525	0.1193	0.1501	0.1672
KF	0.0792	0.1230	0.1632	0.1136	0.1460	0.1625
LSE	**0.0688**	**0.1073**	**0.1398**	**0.1042**	0.1345	0.1501
LSE+VR	0.0695	0.1075	0.1407	0.1043	**0.1338**	**0.1492**
Mask coverage	100%	32%	8%	100%	51%	29%

average, with worst observed values below 0.8 ms. In the jet case, this overhead is comfortably below the 10 ms time budget associated with the 100 Hz update rate, indicating that the reduced-order refinement is not the limiting factor for real-time deployment. In the channel-flow case, the available time budget would be more restrictive because of the higher acquisition frequency. On the other hand, the number of spatial degrees of freedom N_p^{HR} is lower by a few times, so the cost of the reduced-order operations is also expected to be correspondingly smaller.

The timing breakdown also clarifies the computational share of the steps of the process. The projection of the LR field and the latent-state estimation remain comparatively inexpensive, whereas the dominant contribution is the reconstruction of the full HR field. This trend is expected from the dimensionality of the three operations: the LR projection scales as $\mathcal{O}(r\, N_p^{\mathrm{LR}})$, the latent-state estimation as $\mathcal{O}(r^2)$, and the full-field reconstruction as $\mathcal{O}(N_p^{\mathrm{HR}}\, r)$. Since $N_p^{\mathrm{HR}} \gg r$, the reconstruction step naturally dominates the total overhead.

This distinction is particularly relevant for control-oriented applications. When the objective is feedback estimation or low-order state monitoring, the online procedure can stop at the latent HR representation, without reconstructing the full field at each time step. In Fig. 13, this case corresponds to the line labelled *Latent*, which includes only the reduced space operations and excludes full-field reconstruction. Its cost remains of the order of 10 µs, well below the characteristic time required for the LR velocity field estimation itself. The proposed framework is therefore especially attractive when the reduced coordinates, rather than the full reconstructed field, are the quantities of primary interest.

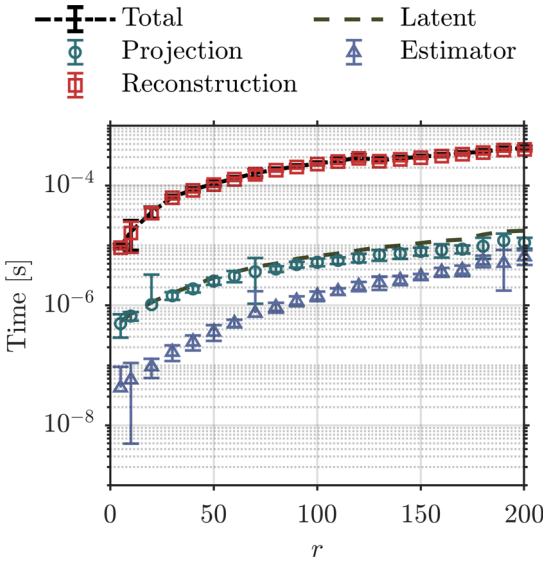


Fig. 13 Computational cost of the additional online reduced-order refinement stage for the jet case as a function of retained rank r. The reported timings distinguish the contribution of the LR projection, latent-state estimation, and HR field reconstruction, together with their total cost. The line labelled *Latent* (i.e. operating only on the reduced-order POD coordinates) denotes the cost of operating only in the reduced space, i.e. projection plus latent-state estimation without full-field reconstruction. This quantity is particularly relevant for control-oriented implementations, in which the latent HR representation is sufficient and the full reconstructed field is not required online. Error bars denote the standard deviation over repeated runs

5 Conclusions

A data-driven reduced-order framework has been proposed to estimate, in real time, high-resolution velocity fields and reduced-order flow coordinates from coarse real-time Event-Based Imaging Velocimetry (rt-EBIV) measurements. The approach was assessed on two broadband turbulent flow datasets, namely a submerged water jet and a channel flow over a square rib, using paired LR/HR data obtained from the same raw EBV recordings through two processing pipelines of different fidelity. As a result, the offline training stage does not require any modification of the experimental setup or additional instrumentation, since the information used to learn the estimator is extracted from the same measurement chain that is employed online.

Compared with direct upsampling with cubic interpolation of the LR measurements, the proposed estimators provide consistently more reliable reconstructions of fluctuation-related flow content. Their benefit is not lim-

ited to visually plausible super-resolution, but extends to a more physically consistent reconstruction of second-order statistics, temporal and spatial spectra, reduced-order modal dynamics, and temporal autocorrelation, while avoiding the grid-imprinting artefacts introduced by purely spatial upsampling. In this sense, the proposed framework should not be interpreted as a simple LR-to-HR interpolation operator, but rather as a reduced-order estimator that combines learned cross-resolution mappings with temporal regularization.

An important aspect of the methodology is that the reduced-order modelling stage also acts as a denoising mechanism. This is particularly relevant in the present context, where the online LR measurements are obtained from a deliberately lightweight single-pass processing chain, which is computationally efficient but more sensitive to noise than the offline multi-pass procedure used to construct the HR reference. The signature of this excess noise is directly visible in the spectral analyses of Sect. 4.4, where the raw LR spectra exhibit an elevated broadband floor at high frequencies and wavenumbers (Figs. 8 and 9). Direct interpolation propagates this noisy and under-resolved information to the finer grid, whereas the proposed estimators constrain the reconstruction to a learned reduced-order subspace. This effect is especially valuable when the reduced-order HR coordinates themselves are the quantities of interest, since operating directly in reduced space provides access to a compact and dynamically meaningful representation of the flow while filtering part of the measurement noise.

Across the different diagnostics, no single estimator dominates uniformly in every representation. LSE yields the lowest global scalar error, LSE+VR remains very close to it and improves the recovery of fluctuation energy and higher-order content, while KF remains robust in several field- and spectrum-based comparisons. The variance rescaling correction should nevertheless be interpreted with care, since it is heuristic and identified from empirical statistics rather than derived from an optimal estimation principle. Its value is therefore mainly pragmatic: it compensates part of the systematic attenuation of standard LSE without changing the overall reduced-order estimation framework.

From a practical standpoint, the additional online cost of the refinement stage remains small compared with the real-time LR measurement itself. In particular, when only the reduced-order HR representation is required, the method can operate entirely in a reduced space at very low cost. The presence of a learned dynamical model also provides an additional practical advantage: if an LR measurement is temporarily unavailable, for example because of frame skipping, deliberate downsampling, or a transient processing bottleneck, the reduced HR state can still be propagated for a limited time using the model prediction alone. Although such a prediction-only estimate is less reliable than a measurement-corrected one, it remains preferable to a complete loss of the state estimate over short horizons.

The present methodology nevertheless remains tied to the operating conditions represented in the training data. In its current form, it should therefore not be regarded as a general solution for actuation-modified or strongly off-design flow regimes. This limitation should, however, be put in perspective. First, the offline training stage is required only once: its role is analogous to that of a calibration procedure, after which the measurement chain delivers offline-grade reduced-order estimates at real-time rates that the offline pipeline could not sustain. The cost of this calibration is quantified in Appendix A: training records of the order of 10^2 convective time units suffice for the estimators to outperform the interpolation baseline, and about 300 convective time units recover the full performance reported in Table 2. This corresponds to less than one second of acquisition for the channel-flow case. The same train-offline/estimate-online paradigm underpins established real-time estimation frameworks such as SPPIV and its closed-loop applications (Kanda et al. 2022; Nonomura et al. 2025; Viguera et al. 2026). Second, the real-time monitoring and diagnostics of unactuated flows (e.g. continuous monitoring of experiments, detection of transients, or virtual sensing in reduced-order coordinates) constitutes a complete application domain in which the method is directly deployable in its present form. The one-off calibration is therefore performed only once before deployment: since tracking the reduced-order state requires storing only a few hundred coefficients per snapshot at negligible computational cost (Sect. 4.8), acquisitions substantially longer than the offline processing could afford become accessible, easing the convergence of statistics that require long records. Third, the framework is already applicable to a relevant class of flow-control problems, namely those in which the actuation does not significantly alter the dynamics within the sensed region. A notable example is the configuration in which sensing, actuation, and performance evaluation are located in distinct areas of the region of interest, so that the actuation affects the reward region but only marginally the sensed one, as in the real-time EBIV control experiment presented by Willert et al. (2026). Extending the framework towards actuation-aware estimation (i.e. towards fully closed-loop active-flow-control conditions in which the actuation modifies the sensed dynamics) and multi-regime training is a natural next step; a promising route, common in the flow-control literature, is to train the estimator over a basis of spanned actuation policies, so that the Kalman filter operates on a reduced-order representation valid across the actuated regimes. In addition, the present formulation relies on a linear reduced-order transition model and on a linear POD-based encoding. Future developments may therefore explore nonlinear extensions of both ingredients, for example through operator-inference strategies and nonlinear filtering

frameworks such as the extended or unscented Kalman filter, which accommodate nonlinear state-transition and measurement models through local linearization and deterministic sampling of the state statistics, respectively, as well as nonlinear reduced-order representations based on autoencoders or related machine learning approaches. Finally, the development of a validated synthetic continuous- and pulsed-EBIV simulator (capable of reproducing the sensor behaviour under pulsed illumination of tracer particles, including per-pixel threshold dispersion, timing jitter, and arbiter saturation) would provide access to ground-truth-based assessments and represents, in the authors' view, a valuable development for the event-based velocimetry community at large.

Within its current scope, however, the results demonstrate that deliberately coarse rt-EBIV processing can be used to extend the effective real-time operating range of the measurement chain towards higher sampling frequencies, while still recovering substantially richer and dynamically more consistent HR flow representations through reduced-order estimation. This is promising not only for high-quality real-time flow diagnostics, but also as a first step towards future observer-based and closed-loop flow-control applications.

Appendix A: Sensitivity to the training-dataset size

The offline training stage of the proposed framework plays the role of a calibration procedure, and its practical cost is set by the length of the recording required to identify the reduced-order bases and operators. This appendix quantifies this cost by repeating the complete training procedure on progressively shorter portions of the training data.

For each training-set size $N_{train} \in \{500, 1000, 1500, 2250, 3000, 4500\}$, the KF and LSE estimators are trained on contiguous windows of snapshots. Five window positions are considered within the full-training range, except for the largest training-set size, for which only one window is available. For each window, the estimators are built using only the information available within that window, namely the mean flow, the POD bases, the truncation ranks selected with the *elbow criterion* described in Sect. 2.1, the mapping operators, and the noise covariances. The validation and testing subsets are kept fixed. Performance is evaluated using the metric of Eq. (34) against the same HR-LOR reference used throughout the paper, which is built from the full-training set. Therefore, $\delta(N_{train})$ quantifies how closely a shorter calibration reproduces the full-training product. The comparison is performed on the total velocity fields, so that errors in the mean-flow estimation from shorter records are also included.

The results are reported in Fig. 14. The error decreases monotonically with the training-set size, with a limited scatter across the window positions, and the two estimators behave very similarly. Both match the full-training cubic-interpolation baseline with training records of about 1500 snapshots for the jet and 2250–3000 for the channel flow: expressed in physical terms, records of the order of 10^2 convective time units in both cases (i.e. 15–30 s of acquisition for the jet and 0.3–0.6 s for the channel flow). Approximately 300 convective time units recover the full-training performance of Table 2, corresponding to 45 s of acquisition for the jet and to less than one second (0.9 s) for the channel flow. At the shortest records considered (≈ 33 convective time

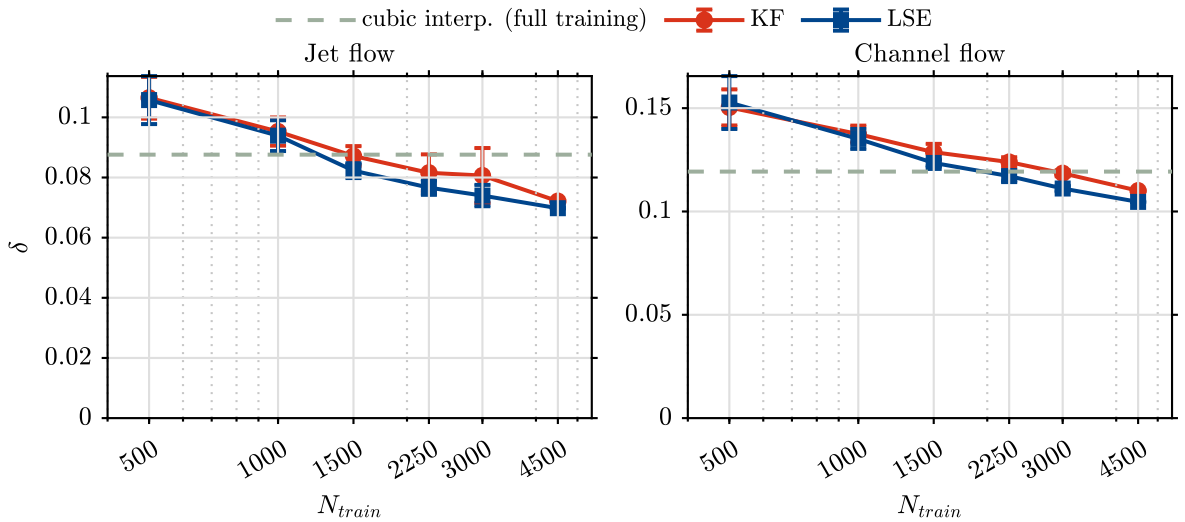


Fig. 14 Global reconstruction error δ (Eq. 34, evaluated against the full-training HR-LOR reference) as a function of the training-dataset size, for the jet (left) and channel-flow (right) cases. Symbols and error bars denote the mean and standard deviation over the five training-window positions. The dashed line reports the full-training cubic-interpolation baseline of Table 2

units), the estimation degrades gracefully, with δ exceeding the interpolation baseline by roughly 20–30%.

Two remarks are due for a fair reading of these curves. First, δ is the most severe possible metric for the short-training estimators, since it combines the pointwise character of the metric discussed in Sect. 4.7—which favours smoother estimates over dynamically consistent ones—with the representational mismatch between the coarser training-window basis and the fixed full-training reference basis; the values at small N_{train} should therefore be regarded as conservative bounds on the practical usability. Second, the truncation rank selected by the *elbow criterion* can oscillate significantly between training windows, with only a marginal impact on δ, indicating that the performance of the framework is robust with respect to the exact truncation choice. Overall, these results support the interpretation of the offline training stage as a lightweight, one-off calibration: a recording of the order of 10^2 convective time units suffices to deploy estimators that outperform direct interpolation, with performance saturating towards the full-training level as the calibration record grows.

Acknowledgements This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 949085, NEXTFLOW ERC StG). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The help of Michael Schroll of the DLR Institute of Propulsion Technology was invaluable in setting up the air flow experiment.

Author contributions L.F. helped in conceptualization, data curation, formal analysis, investigation, methodology, software, writing—original draft, writing—review and editing. E.A. contributed to formal analysis, investigation, methodology, software, writing—review and editing. C.E.W. was involved in investigation, resources, software, writing—review and editing. M.R. helped in methodology, supervision, writing—review and editing. G.C. contributed to supervision, writing—review and editing. S.D. helped in conceptualization, formal analysis, funding acquisition, methodology, project administration, resources, supervision, writing—review and editing.

Data availability Data supporting this work are available on Zenodo at https://zenodo.org/uploads/20037404. The archive includes the `.raw` event-camera recordings, a subset of the corresponding pseudo-snapshots, and the processed velocity fields used in the present analysis. Additional data are available from the corresponding author upon reasonable request. A Python implementation of the real-time EBIV processing workflow and of the high-resolution estimation framework proposed in this work is available at https://github.com/LFranceschelli/VIBEstream. This repository is intended to document and reproduce the estimation methodology; it does not include the full C++ implementation used for the dedicated real-time benchmark.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

AI and AI-assisted technologies in the writing process During the preparation of this work, the authors used ChatGPT and Grammarly to improve the readability and language of this manuscript. After using this tool/service, the authors reviewed and edited the content as needed and took full responsibility for the publication's content.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Amico E, Zannone M, Torta E, Mastronuzzi G, Pecchio D, Gallo D, Serpieri J, Cafiero G, Morbiducci U (2026) Event-based versus particle image velocimetry for cardiac flow analysis in a left-heart simulator. Phys Fluids 38(2):021916. https://doi.org/10.1063/5.0315234

Astarita T (2007) Analysis of weighting windows for image deformation methods in PIV. Exp Fluids 43(6):859–872

Audiffred DBS, Cavalieri AVG, Maia IA, Martini E, Jordan P (2024) Reactive experimental control of turbulent jets. J Fluid Mech 994:15. https://doi.org/10.1017/jfm.2024.569

Berkooz G, Holmes P, Lumley J (2003) The proper orthogonal decomposition in the analysis of turbulent flows. Annu Rev Fluid Mech 25:539–575. https://doi.org/10.1146/annurev.fl.25.010193.002543

Bollt SA, Foxman SH, Gharib M (2025) RapidPIV: full flow-field kHz PIV for real-time display and control. arXiv preprint arXiv:2504.17987, https://doi.org/10.48550/arXiv.2504.17987

Brindise MC, Vlachos PP (2017) Proper orthogonal decomposition truncation method for data denoising and order reduction. Exp Fluids 58(4):28. https://doi.org/10.1007/s00348-017-2320-3

Cortina-Fernández J, Sanmiguel Vila C, Ianiro A, Discetti S (2021) From sparse data to high-resolution fields: ensemble particle modes as a basis for high-resolution flow characterization. Exp Thermal Fluid Sci 120:110178. https://doi.org/10.1016/j.expthermflusci.2020.110178

Discetti S, Bellani G, Örlü R, Serpieri J, Sanmiguel Vila C, Raiola M, Zheng X, Mascotelli L, Talamelli A, Ianiro A (2019) Characterization of very-large-scale motions in high-Re pipe flows. Exp Thermal Fluid Sci 104:1–8. https://doi.org/10.1016/j.expthermflusci.2019.02.001

Dacome G, Mörsch R, Kotsonis M, Baars WJ (2024) Opposition flow control for reducing skin-friction drag of a turbulent boundary layer. Phys Rev Fluids 9(6):064602. https://doi.org/10.1103/PhysRevFluids.9.064602

Deng Z, He C, Liu Y, Kim KC (2019) Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework. Phys Fluids 31(12):125111. https://doi.org/10.1063/1.5127031

Drazen D, Lichtsteiner P, Häfliger P, Delbrück T, Jensen A (2011) Toward real-time particle tracking using an event-based dynamic vision sensor. Exp Fluids 51:1465–1469. https://doi.org/10.1007/s00348-011-1207-y

Epps BP, Krivitzky EM (2019) Singular value decomposition of noisy data: mode corruption. Exp Fluids 60(8):121. https://doi.org/10.1007/s00348-019-2761-y

Epps BP, Techet AH (2010) An error threshold criterion for singular value decomposition modes extracted from PIV data. Exp Fluids 48(2):355–367. https://doi.org/10.1007/s00348-009-0740-4

Ewing D, Citriniti JH (1999) Examination of a LSE/POD complementary technique using single and multi-time information in the axisymmetric shear layer. In: Sørensen JN, Hopfinger EJ, Aubry N (eds.) IUTAM symposium on simulation and identification of organized structures in flows. Springer, Dordrecht, pp 375–384. https://doi.org/10.1007/978-94-011-4601-2_33

Foucaut J-M, Carlier J, Stanislas M (2004) PIV optimization for the study of turbulent flow using spectral analysis. Meas Sci Technol 15(6):1046–1058. https://doi.org/10.1088/0957-0233/15/6/003

Franceschelli L, Willert CE, Raiola M, Discetti S (2025) An assessment of event-based imaging velocimetry for efficient estimation of low-dimensional coordinates in turbulent flows. Exp Thermal Fluid Sci 164:111425. https://doi.org/10.1016/j.expthermflusci.2025.111425

Fukami K, Fukagata K, Taira K (2021) Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. J Fluid Mech 909:9. https://doi.org/10.1017/jfm.2020.948

Gallego G, Delbrück T, Orchard G, Bartolozzi C, Taba B, Censi A, Leutenegger S, Davison AJ, Conradt J, Daniilidis K, Scaramuzza D (2022) Event-based vision: a survey. IEEE Trans Pattern Anal Mach Intell 44(1):154–180. https://doi.org/10.1109/TPAMI.2020.3008413

Gautier N, Aider J-L, Duriez T, Noack BR, Segond M, Abel M (2015) Closed-loop separation control using machine learning. J Fluid Mech 770:442–457. https://doi.org/10.1017/jfm.2015.95

Glauser MN, George WK (1987) Orthogonal decomposition of the axisymmetric jet mixing layer including azimuthal dependence. In: Comte-Bellot G, Mathieu J (eds.) Advances in turbulence. Springer, Berlin, Heidelberg, pp 357–366. https://doi.org/10.1007/978-3-642-83045-7_40

Howell J, C Hammarton T, Altmann Y, Jimenez M (2020) High-speed particle detection and tracking in microfluidic devices using event-based sensing. Lab Chip 20(16):3024–3035. https://doi.org/10.1039/D0LC00556H

Kalman RE (1960) A new approach to linear filtering and prediction problems. J Basic Eng 82(1):35–45. https://doi.org/10.1115/1.3662552

Kanda N, Abe C, Goto S, Yamada K, Nakai K, Saito Y, Asai K, Nonomura T (2022) Proof-of-concept study of sparse processing particle image velocimetry for real time flow observation. Exp Fluids 63(9):143. https://doi.org/10.1007/s00348-022-03471-0

Kreizer M, Ratner D, Liberzon A (2010) Real-time image processing for particle tracking velocimetry. Exp Fluids 48:105–110. https://doi.org/10.1007/s00348-009-0715-5

Lichtsteiner P, Posch C, Delbruck T (2008) A 128 × 128 120 dB 15 μs latency asynchronous temporal contrast vision sensor. IEEE J Solid-State Circuits 43(2):566–576. https://doi.org/10.1109/JSSC.2007.914337

Mahowald MA, Mead C (1991) The silicon retina. Sci Am 264(5):76–82. https://doi.org/10.1038/scientificamerican0591-76

McCormick F, Gibeau B, Ghaemi S (2024) Reactive control of velocity fluctuations using an active deformable surface and real-time PIV. J Fluid Mech 985:9. https://doi.org/10.1017/jfm.2024.292

Nonomura T, Abe C, Naramura R, Sasaki Y (2025) Real-time feedback control of flow velocity field using sparse processing particle image velocimetry and plasma actuators. Exp Fluids 66(7):136. https://doi.org/10.1007/s00348-025-04039-4

Pimienta J, Aider J-L (2025) High resolution and high-speed live optical flow velocimetry. https://arxiv.org/abs/2509.25924v3

Raffel M, Willert CE, Kähler CJ, Scarano F, Wereley ST, Kompenhans J (2018) Particle image velocimetry: a practical guide. 3rd Ed. Springer, Berlin Heidelberg. https://doi.org/10.1007/978-3-319-68852-7

Raiola M, Discetti S, Ianiro A (2015) On PIV random error minimization with optimal POD-based low-order reconstruction. Exp Fluids 56(4):75. https://doi.org/10.1007/s00348-015-1940-8

Rajaee M, Karlsson SKF, Sirovich L (1994) Low-dimensional description of free-shear-flow coherent structures and their dynamical behaviour. J Fluid Mech 258:1–29. https://doi.org/10.1017/S0022112094003228

Rusch A (2023) TrackAER: volumetric event-based particle tracking velocimetry in real-time. PhD thesis, ETH Zürich. https://doi.org/10.3929/ethz-b-000603611

Rusch A, Rösgen T (2023) TrackAER: real-time event-based quantitative flow visualization. Exp Fluids 64:136. https://doi.org/10.1007/s00348-023-03673-0

Scarano F (2002) Iterative image deformation methods in PIV. Meas Sci Technol 13(1):1–19. https://doi.org/10.1088/0957-0233/13/1/201

Shiba S, Hamann F, Aoki Y, Gallego G (2024) Event-based background-oriented schlieren. IEEE Trans Pattern Anal Mach Intell 46(4):2011–2026. https://doi.org/10.1109/TPAMI.2023.3328188

Siegel S, Cohen K, McLaughlin T, Myatt J (2003) Real-time particle image velocimetry for closed-loop flow control studies. In: 41st aerospace sciences meeting and exhibit, Reno (NV). https://doi.org/10.2514/6.2003-920

Sirovich L (1987) Turbulence and the dynamics of coherent structures. I Coherent structures Q Appl Math 45(3):561–571. https://doi.org/10.1090/qam/910462

Stankovic JA (2002) Misconceptions about real-time computing: a serious problem for next-generation systems. Computer 21(10):10–19

Tinney CE, Coiffet F, Delville J, Hall AM, Jordan P, Glauser MN (2006) On spectral linear stochastic estimation. Exp Fluids 41(5):763–775. https://doi.org/10.1007/s00348-006-0199-5

Tirelli I, Ianiro A, Discetti S (2023) An end-to-end KNN-based PTV approach for high-resolution measurements and uncertainty quantification. Exp Thermal Fluid Sci 140:110756. https://doi.org/10.1016/j.expthermflusci.2022.110756

Tu JH, Griffin J, Hart A, Rowley CW, Cattafesta LN, Ukeiley LS (2013) Integration of non-time-resolved PIV and time-resolved velocity point sensors for dynamic estimation of velocity fields. Exp Fluids 54(2):1429. https://doi.org/10.1007/s00348-012-1429-7

Varon E, Aider J-L, Eulalie Y, Edwige S, Gilotte P (2019) Adaptive control of the dynamics of a fully turbulent bimodal wake using real-time PIV. Exp Fluids. https://doi.org/10.1007/s00348-019-2766-6

Viguera R, Naramura R, Sasaki Y, Nonomura T (2026) Adaptive control pattern for real-time-visual-feedback flow separation control over airfoil with sparse processing particle image velocimetry and plasma actuator. Exp Fluids. https://doi.org/10.1007/s00348-026-04175-5

Welch G, Bishop G (1995) An introduction to the Kalman filter. https://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf

Westerweel J, Scarano F (2005) Universal outlier detection for PIV data. Exp Fluids 39(6):1096–1100. https://doi.org/10.1007/s00348-005-0016-6

Willert C (2023) Event-based imaging velocimetry using pulsed illumination. Exp Fluids 64:98. https://doi.org/10.1007/s00348-023-03641-8

Willert C (2025) Event-based particle image velocimetry for high-speed flows. Meas Sci Technol 36(7):075302. https://doi.org/10.1088/1361-6501/ade27a

Willert C, Munson M, Gharib M (2010) Real-time particle image velocimetry for closed-loop flow control applications. In: 15th international symposium on applications of laser techniques to fluid mechanics

Willert C, Franceschelli L, Amico E, Raiola M, Cafiero G, Discetti S (2026) Real-time event-based particle image velocimetry for active flow control. J Phys: Conf Ser 3173(1):012001. https://doi.org/10.1088/1742-6596/3173/1/012001

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.